Recherches sur la théorie de la décision

Dossier de travaux présenté en vue de l'obtention du

Diplôme d'Habilitation à Diriger des Recherches
en
Sciences Économiques
de
l'Université Paris 1 Panthéon-Sorbonne

par

Brian Hill

Chargé de recherche au CNRS,
et Professeur affilié à HEC Paris.

Note pour le lecteur / Note for the Reader

Ce document est composé d'une introduction, ou *Note de Synthèse* (en français), et neuf articles (en anglais). Chaque "chapitre" a ses propres annexe et références, ainsi qu'une numérotation des pages, sections, théorèmes et ainsi de suite qui lui est spécifique. Pour cette raison, chaque page, section, et ainsi de suite comporte le numéro de l'article : par exemple, §V.2.2 est la section 2.2 de l'article V.

Malgré sa maladresse, cette numérotation est utilisée avec l'intention de faciliter la lecture de et la référence au document. Même s'il y a un ordre naturel, exposé dans la *Note de Synthèse*, les articles peuvent être lus indépendamment les uns des autres, et dans n'importe quel ordre. Ainsi, pour le lecteur qui souhaite profiter de cette possibilité, une numérotation des théorèmes, axiomes et ainsi de suite qui indique le contexte dans lequel ils sont discutés nous a paru plus commode.

The document before you is composed of an Introduction, or *Note de Synthèse* (in French), and nine papers (in English). Each of these 'chapters' is entirely self-sufficient, with its own appendices and bibliography, and its own numbering of pages, sections, theorems, axioms and so on. To ensure this, each page, section and so on is prefixed with the paper number : for example, §V.2.2 is subsection 2 of section 2 from Paper V, Axiom VII-A4 is the A4 Axiom in Paper VII, page III-13 is page 13 of Paper III.

Despite its clumsiness, this numbering has been adopted to facilitate reading and reference to the document. Although there is a sort of natural order, explained in the *Note de Synthèse*, each paper (and the Introduction) can be read independently of the others, and they can be read in any order, if wished. Moreover, to aid quick and easy reference to axioms, theorems and so on, it seemed easier to prefix them by their paper number to indicate the context in which they are discussed.

Acknowledgements

This document collects some of my research in decision theory spanning almost eight years, and its preparation has given me the opportunity to look back over this period. A lot has happened to me in this time, and I have much to be grateful for.

I am grateful, first of all, to a community. An HDR is a 'sort of second thesis', but this is my first thesis in decision theory. This in itself is a testament to the open-mindedness, patience and kindness of researchers in this domain, without which many of the papers presented below would not exist. Coming from a different background, I could not help being surprised by the vibrancy and humanity of the decision theory community, internationally but also, and in particular, in France. Researchers working in France have had a special role to play in the construction of the field – in all senses of the term – and I consider myself particularly fortunate to now belong to this community, and to have had the opportunity to interact with many of its members. A great thanks, then, to conference participants, colleagues, and co-authors who I have had the pleasure and privilege to discuss with, to get feedback from, and to learn from, and who are too numerous to list here. I have, of course, a particular debt of gratitude to Alain Chateauneuf, Michèle Cohen, Edi Karni, and Jean-Marc Tallon who, beyond the encouragement and intellectual stimulation which they have given me over the years, have accepted to participate in this jury, taking on the responsibility of *rapporteur* in the case of Alain Chateauneuf and Edi Karni, and the role of *directeur*, with remarkable patience and availability, in the case of Jean-Marc Tallon.

I am grateful, secondly, to a research group. The Decision Sciences group at GREGHEC, beyond providing excellent research assistance, has furnished a superb intellectual environment. To all those involved in this environment, and in creating it, I owe a debt of gratitude. Particular thanks is due to Nathalie Beauchamp, whose efficiency and good cheer has rendered my professional life not only easier but more pleasurable, and to those in the group and associated with it with whom I have had most interactions, Stefania Minardi, Emmanuel Kemel, Antoine Nebout (from whom I also appreciate the effort and patience required to check my French) and Tomasz Michalski. I am especially grateful to Mohammed Abdellaoui for his valuable advice and interesting discussions, and to Itzhak Gilboa, who, beyond the priceless support and stimulating exchanges, has been a role model of intellectual curiosity and integrity. He also deserves a special thanks for his availability, exemplified by his participation in this jury as *rapporteur*, despite the various constraints on his agenda.

I am grateful, also, to the person who introduced me to decision theory. It is no accident that Philippe Mongin is alone in having had the pleasure of sitting on both my thesis and HDR juries (and of having read both) – an achievement that in itself merits recognition. He more than anyone is ultimately responsable for the fact that this HDR is in decision theory. He has not only managed to communicate his enthusiasm for many of the questions discussed in the following pages, but has taught me to appreciate their subtlety and inspired me to tackle them. For this, and for his unfailing support and encouragement, I am profoundly grateful.

The topic of decision-making is relevant in many walks of life, and anyone's thinking on such issues will doubtless be influenced by the interactions and experiences he has outside the narrow field of decision theory. I am grateful to all working beyond economics who I have had the good fortune to discuss with and learn from, and in particular Richard Bradley, whose enthusiasm, curiosity and openness have fuelled many a thought-provoking discussion. I also consider myself lucky to have had frequent fruitful exchanges with members of the DRI research group, and in particular Mikaël Cozic.

But any influence stretches beyond even the academic environment, and hence I must mention my appreciation of the many acquaintances and especially friends (some of which are mentioned above in their role as researchers) who have offered their wisdom, their opinions and the opportunity to talk over these and other topics. To these friends, and to my family and extended family, in particular Dave, Ian, Katia, Cesare, Marco and Sonia, I also owe a huge thanks for their affection, their humour and their support. Last but not least, for all the reasons mentioned above, and so many more, I am infinitely grateful to Francesca and to the most important thing that has happened in the last eight years, Vincent : the easiest decisions of my life, without whom I could not have faced the toughest ones.

Paris, January 2015

Table des matières

I Note de synthèse

I.1 Introduction

I.1.1 La décision comme recette

Imaginez que vous deviez choisir le plan pluriannuel de politique environnementale pour un pays. Prendre une décision comme celle-ci est, à certains égards, comme cuisiner un plat. Il y a des ingrédients : vos propres désirs et goûts et ceux des gens que vous représentez en tant que décideur, mais aussi vos croyances, informées dans la mesure du possible par les meilleures recherches scientifiques. Il y a un produit : la décision prise. Et il y a des manières de traiter, manipuler, combiner les premiers pour produire le dernier : la recette, si vous voulez. Une des préoccupations de la théorie de la décision, sinon sa préoccupation majeure, est de comprendre ce passage des ingrédients à la décision, de comprendre la recette de la décision.

On peut s'intéresser à ce passage à trois titres au moins. Tel le chef qui cherche la bonne recette, on peut vouloir savoir quelle serait la meilleure façon de prendre des décisions. En principe, la théorie de la décision peut fournir des réponses à cette question *normative*, sous forme des recettes qui sont "bonnes", accompagnées en tant que possible d'une explication ou défense de leurs qualités. Par ailleurs, on peut vouloir savoir quelles recettes les cuisiniers ordinaires utilisent réellement. Plutôt que le chef qui veut cuisiner le meilleur plat, on se met dans le rôle de l'anthropologue ou du psychologue, qui, sur la base des plats produits (les décisions prises) veulent comprendre les recettes sous-jaçentes. Des théoriciens de la décision affichent souvent l'ambition de fournir des éléments de réponse à cette question *descriptive*. Enfin, même si les cuisiniers ordinaires utilisent une recette très compliquée, une recette facile qui produise en fait plus ou moins la même chose est parfois désirable.

L'économiste qui étudie l'assurance, les marchés financiers, ou l'organisation industrielle réclame un modèle disant comment les acteurs économiques prennent des décisions. Pour ses objectifs, il est prêt à sacrifier l'exactitude à la simplicité, et la théorie de la décision est sensible à ce besoin. Pour ce qui concerne l'économie au moins, l'importance première du domaine provient sans doute de cette dimension *de modélisation* ou *d'application économiques*.

Une recette – un modèle de décision, ou une théorie comme on l'appelle parfois – ne répond pas à elle seule à ces différentes questions. Il en faut aussi une défense, voire une explication. D'un côté, on aimerait pouvoir *évaluer* le modèle, en particulier sur la base de ses propriétés comportementales, et se servir d'une telle évaluation pour juger sa plausibilité en tant que réponse à l'une ou l'autre des trois questions. D'un autre côté, on voudrait que le modèle et ses composantes aient un *fondement*, notamment dans les comportements observables. Il est souvent important de savoir comment *identifier* les différentes composantes du modèle : comment mesurer leurs valeurs chez des décideurs particuliers, par exemple. L'usage correct du modèle dans les applications économiques présuppose également une interprétation suffisamment claire de ces composantes, et une compréhension de leurs propriétés.

Dans le champ de la décision dans l'incertain – où, comme pour la décision environnementale ci-dessus, les probabilités des résultats ne sont pas données de manière exogène – les théoriciens de la décision ont prétendu fournir des réponses aux trois questions précédentes, ainsi que des moyens pour évaluer leurs propositions. D'ailleurs, pendant un certain nombre d'années, tout se passait comme s'il y avait la même réponse chaque fois : la *théorie de l'espérance d'utilité*. Elle est importante non seulement en tant que point de départ pour toute théorie concurrente, mais également en raison des outils d'évaluation et d'identification qu'on a développés pour elle, qui font toujours référence dans le domaine. Comme pour toute présentation de la théorie de la décision dans l'incertitude, notre point de départ est donc la théorie de l'espérance d'utilité.

I.1.2 L'espérance de l'utilité et ses atouts

Supposons que les objets de choix soient des *actes* : des options qui, pour chaque état de la nature, donnent une conséquence particulière. Un pari hippique, par exemple, est une option de ce genre – selon l'état de la nature (et en particulier le cheval qui gagne la course), il y aura des conséquences, monétaires en l'occurence, différentes. Plus formellement, si l'on dénote l'ensemble des conséquences possibles par C et l'ensemble des états de la nature par S, alors un acte est une fonction de S à C. Comme nous enseigne la théorie du choix, sous certaines hypothèses on peut représenter les choix de l'agent par une relation de préférence faible $\preceq$ sur les actes. Quand il s'agit de décision dans l'incertain, on prend souvent cette relation de préférence comme élément primitif, en supposant qu'il résume le comportement du décideur. Cette convention est adoptée ici sans discussion.

La recette standard – *la théorie de l'espérance d'utilité*, qui va parfois aussi sous le nom de *Bayésianisme* – a deux ingrédients. Le premier est une fonction (réelle) sur les conséquences, qu'on appelle *fonction d'utilité*. On considère souvent qu'elle représente les *désirs* du décideur relatifs aux conséquences. Le second est une *mesure de probabilité* sur les états possibles de la nature. Elle aussi correspondrait à certains états mentaux du décideur, à savoir ses *croyances*. Notons la fonction d'utilité par u et la mesure de probabilité par p. La recette est simple : préférer l'acte dont l'espérance de l'utilité de la conséquence résultante, calculée avec la mesure de probabilité, est la plus grande. En d'autres termes, la préférence est représentée par une fonction V sur les actes – en ce sens standard où $f \preceq g$ si et seulement si $V(f) \leqslant V(g)$ – qui, pour un acte f, est définie de manière suivante : [1]

$$V(f) = \sum_{s \in S} p(s)u(f(s)) \tag{I.1}$$

La pierre angulaire de la théorie de l'espérance d'utilité est un résultat mathématique (voire une famille de résultats) qu'on appelle *théorème de représentation*. Établi dans les travaux novateurs de Ramsey (1931), de Savage (1954), mais aussi d'Anscombe and Aumann (1963), ce résultat fournit un ensemble de conditions sur les préférences – les *axiomes* – et montre une équivalence entre la satisfaction de ces axiomes par une relation de préférence et l'existence d'une fonction d'utilité et d'une mesure de probabilité qui les représentent selon (I.1). De plus, il affirme que cette fonction et cette mesure sont uniques en un sens approprié. En d'autres termes, le résultat présente les conséquences comportementales d'utiliser la recette d'utilité espérée, mais permet également de repérer quand un plat (un ensemble de décisions) peut avoir été fait selon cette recette et même d'identifier les éventuels ingrédients (la fonction d'utilité et la mesure de probabilité).

1. Pour les besoins de cette présentation, on peut supposer que l'espace d'états S est fini.

L'importance du théorème de représentation vient du fait qu'il fournit des outils d'évaluation et d'identification qui, comme nous l'avons constaté ci-dessus, sont très appréciés dans un modèle de décision. D'un côté, puisqu'il repère précisément les propriétés comportementales qui correspondent à la théorie de l'espérance d'utilité, le théorème contribue à l'évaluation de cette théorie. Si l'on trouve que les gens satisfont les axiomes, ou qu'ils les satisfont à peu près en général dans un certain type de situation, alors le modèle d'utilité espérée, qui leur est équivalente, est validée pour des fins descriptives ou de modélisation économique. Si, par contre, un axiome est violé de manière régulière, c'est une indication que le modèle ne réussit pas à représenter le comportement. On peut répéter le même discours *mutatis mutandis* pour la dimension normative. De l'autre côté, par sa clause d'unicité, le théorème se prononce sur l'identification des composantes du modèle, en l'occurence l'utilité et la probabilité. Il indique ce qu'il faut mesurer si l'on veut reconstruire les préférences du décideur, et dans certains cas comment le mesurer. D'après certains, il donne ainsi un sens, et même un *fondement* aux concepts de désir et de croyance que l'on peut relier à l'utilité et à la probabilité (Cozic and Hill, forthcoming). Dans les applications, il suffit de se référer au théorème pour déterminer s'il y a sens à faire varier tel ou tel aspect de l'utilité ou de la probabilité, ou si cette variation ne correspond à rien dans les comportements.

La théorie de l'espérance d'utilité fut le seul plat servi pendant presque un demi-siècle ; pourtant, elle a subi depuis quelques dizaines d'années des objections importantes. Celles-ci ne touchent guère aux aspects fondationnels. En effet, la théorie de l'utilité espérée continue à donner l'exemple en la matière, à ce point que, pour une nouvelle théorie, le développement d'un théorème de représentation, avec une clause d'unicité qui serait idéalement aussi forte que celle de l'utilité espérée, est quasi obligatoire dans le domaine. Les objections à la théorie surviennent plutôt sur le plan de l'évaluation. La recette de l'utilité espérée est la réponse à toutes les questions évoquées ci-dessus, à condition qu'on accepte les axiomes. Pourtant, on sait depuis longtemps que certains axiomes sont violés par les individus dans la pratique. De plus, ces violations semblent parfois tellement régulières qu'on a commencé à renoncer à l'usage de la théorie au profit d'autres pour certaines applications économiques (il y aura des exemples par la suite). Enfin, il n'est pas clair que toutes les violations soient à condamner sur le plan normatif ; tant et si bien que certains ont mis en cause la prétention normative de la théorie de l'espérance d'utilité. Que ce soit donc sur les plans normatif, descriptif ou de la modélisation, des voix s'élèvent contre le standard, et la recherche est lancée pour trouver d'autres recettes.

Un avantage des théorèmes de représentation est qu'ils structurent le débat pour ou contre une théorie autour de ses propriétés comportementales, inscrites dans les axiomes. Dans cette introduction, j'ai préféré une autre organisation, par les aspects de la représentation qui doivent être modifiés si l'objection est acceptée. Dans la gastronomie, une modification de la recette pourrait entraîner un changement dans les ingrédients. La même chose vaut pour la décision : les objections, qui peuvent être formulées comme des violations des axiomes, impliquent de modifier l'un ou l'autre des composantes du modèle (l'utilité ou la probabilité, en l'occurence). Une organisation sur cette base permet de rassembler des modèles avec une communauté profonde, alors que leurs divergences par rapport à l'utilité espérée portent sur des axiomes différents. Organisées de cette manière, les recherches présentées dans la suite s'occupent de deux critiques orthogonales de la théorie de l'espérance d'utilité. Je présente brièvement la motivation et les enjeux de chacune, avant de résumer mes contributions dans §§I.2, I.3 et I.4.

I.1.3 L'utilité dépendante des états

L'objection La première objection contre la représentation par maximisation d'espérance d'utilité tient à l'indépendance supposée de l'utilité par rapport aux états. Comme il est évident par (I.1), l'utilité d'une conséquence dépend d'elle seule et d'aucune manière de l'état de la nature réalisé. Or, comme l'ont fait remarquer plusieurs théoriciens (dont par exemple, Drèze et Aumann dans Drèze (1987), Arrow (1974) et Edi Karni dans plusieurs articles qui seront discutés par la suite), cette hypothèse n'est pas toujours valide. Un exemple banal où elle semble être violée concerne les décisions de santé : €100 000 n'a pas la même valeur pour le décideur quand il est en bonne santé que quand il est malade, handicapé ou mort. Prenant les €100 000 comme une conséquence et l'état de santé du décideur comme un aspect de l'état de la nature, son utilité (pour les €100 000) est dépendante de l'état (de santé). La littérature sur la *dépendance de l'utilité à l'égard des états* (*state-dependent utility*) trouve son application la plus reconnue dans ce champ : après tout, dans les décisions concernant l'assurance-maladie ou l'assurance-vie, il faut peser la valeur de l'argent dans un état de santé contre sa valeur dans un état de maladie ou de mort. Il convient de souligner cependant que la question est d'importance même dans les domaines où elle est moins reconnue. Un exemple est fourni par les décisions concernant la politique environnementale mentionnées au tout début. La valeur de l'argent, et d'ailleurs de plusieurs biens, est sans doute différente dans un climat normal et dans des situations de chaleur ou de froid extrême, d'aridité, de déluges ou d'inondation. Comme dans le cas de la santé, cette dépendance devrait être prise en compte dans les décisions prises *ex ante* en présence d'incertitude concernant les facteurs climatiques.

Les propriétés comportementales La formulation précise des axiomes mis en cause dans ces cas de dépendance de l'utilité à l'égard des états dépend du cadre adopté. Deux cadres formels sont souvent pris comme repères dans la littérature. Celui de Savage, proposé dans son ouvrage marquant (Savage, 1954), ne fait d'hypothèse ni sur la nature des conséquences ni sur l'espace d'états.[2] Celui d'Anscombe-Aumann, développé quelques années plus tard (Anscombe and Aumann, 1963) et devenu plus populaire grâce à sa facilité technique plus grande, suppose d'emblée que les conséquences sont des loteries au sens de von Neumann et Morgenstern (des distributions de probabilité à support fini sur un ensemble de résultats).[3] Voir par exemple l'article II pour plus de détails.

Chez Savage, deux axiomes ont trait à l'indépendance de l'utilité à l'égard des états (voir l'article II pour leurs formulations et les détails techniques). L'un – souvent appelé *monotonie* – garantit l'indépendance *ordinale* de l'utilité à l'égard des états ; c'est-à-dire il assure que l'ordre sur les conséquences engendré par la fonction d'utilité est indépendant de l'état (si l'utilité d'une conséquence x est supérieure à celle de y dans un état, il en va ainsi dans tous les états non-nuls). Alors que cette propriété semble naturelle pour certains types de conséquences, telles les conséquences monétaires, il y a des cas où elle n'est pas raisonnable. On peut préférer vivre dans une maison de campagne plutôt qu'en ville quand on est en bonne santé, la préférence étant renversée quand on est en mauvaise santé (pour des raisons de proximité des services de santé, par exemple). L'autre axiome, appelé souvent *probabilité comparative faible*, implique en plus l'indépendance cardinale de l'utilité à l'égard des états : il fait que non seulement l'ordre sur les conséquences, mais aussi les valeurs d'utilité qui leur sont assignées sont indépendantes de l'état. Pour reprendre l'exemple

2. Par contre, ses axiomes, et notamment P6, impliquent une certaine richesse de l'espace d'états.

3. Dans plusieurs travaux récents, on suppose simplement que l'espace des conséquences est un sous-espace convexe d'un espace vectoriel, le cadre d'Anscombe et Aumann pourant être considéré comme un cas spécial. Pour cette raison, je continue où il est besoin de parler de cadre formel d'Anscombe-Aumann dans ces cas.

des conséquences monétaires, sans l'indépendance cardinale, l'attitude au risque (donnée, comme on le sait depuis les travaux de Pratt (1964) et d'Arrow (1971) par la courbure de la fonction d'utilité) peut différer selon l'état de la nature ; avec l'indépendance cardinale, elle est indépendante de lui. Dans le cadre formel d'Anscombe-Aumann, un seul axiome – appelé malheureusement *monotonie* aussi – garantit l'indépendance ordinale et cardinale de l'utilité à l'égard des états. L'article II montre comment séparer la partie concernant l'indépendance ordinale de la partie cardinale.

Le défi Si l'espérance d'utilité ne convient pas dans les circonstances où l'indépendance de l'utilité à l'égard des états est mise en question, alors quelle représentation devrait-on adopter ? À la différence de la littérature sur l'ambiguïté qui sera discutée par la suite, le défi posé ici n'est pas de trouver la bonne règle de décision. Tous les théoriciens qui se penchent sur la question sont d'accord, pour ce qui concerne la question de la dépendance à l'égard des états, qu'une représentation comme la suivante ferait bien l'affaire (sur les plans normatif, descriptif et de la modélisation) :

$$V(f) = \sum_{s \in S} p(s) u(s, f(s)) \tag{I.2}$$

En d'autres termes, il suffit de rendre la fonction d'utilité dépendante des états. Le défi est plutôt de nature fondationnelle : de proposer un théorème de représentation où la probabilité p et l'utilité u sont fournies ou identifiées de manière (suffisamment) unique. Comme il apparaîtra ci-dessous, il ne suffira pas d'enlever les axiomes qui viennent d'être évoqués pour faire cela ; il faudra les remplacer par quelque chose d'autre.

I.1.4 Croyances et probabilités précises

L'objection La deuxième objection à la représentation des préférences par la maximisation de l'espérance d'utilité concerne la mesure de probabilité. Comme il a été dit précédemment, cette fonction peut être interprétée en termes des croyances du décideur ; d'ailleurs, l'une des thèses derrière l'espérance d'utilité, souvent explicitée quand on traite de l'approche dite bayésienne, est que les croyances peuvent être représentées par des probabilités. Or, il y a des raisons de croire que cette thèse est fausse dans certaines situations.

En référence à Knight (1921), qui a introduit une distinction entre le risque – où l'incertitude peut être représentée ou mesurée par une probabilité – et l'incertitude proprement dite – où ce n'est pas le cas –, on qualifie parfois ces situations d'*incertitude knightienne*.[4]

À première vue, l'incertitude knightienne est partout. Il paraît difficile, sinon impossible, de former des croyances ou des jugements probabilistes d'une certaine précision concernant les cours des marchés (après une crise, par exemple), leurs propriétés sous-jacentes (les vraies valeurs des sociétés, par exemple), la politique future (concernant les guerres, par exemple). En décidant d'une politique climatique et énergétique, nous luttons contre notre ignorance des effets de nos actes sur les variables climatiques futures, malgré l'état avancé de la science sur la question. On peut aussi citer les incertitudes concernant les OGM, les nanotechnologies et d'autres sujets à propos desquels on parle de "Principe de Précaution", comme étant des cas vraisemblables d'incertitude knightienne.

Bien sûr, ces intuitions ne suffisent pas à mettre dans l'embarras l'espérance d'utilité, car il se peut que, dans de telles situations, on décide tout de même comme si l'on avait pour croyances des probabilités précises. Pour ce faire, il faut, comme il est de tradition dans la littérature, un phénomène comportemental qu'on puisse relier à ce genre de situations et qui soit logiquement incompatible avec la maximisation de l'espérance d'utilité. Ellsberg (1961) a proposé des contre-exemples qui satisfont ces critères et ils sont souvent pris comme motivation pour explorer des modèles de décision qui affaiblissent les axiomes de l'espérance d'utilité (voir par exemple l'article IX pour une brève présentation et discussion de l'un d'entre eux).

Les propriétés comportementales Comme il en allait pour la dépendance de l'utilité à l'égard des états, la formulation précise de l'axiome récusé dépend du cadre formel utilisé. Dans celui de Savage, et selon la lecture habituelle des exemples d'Ellsberg, il s'agit du principe de la chose sûre (P2 dans la numération de Savage (1954) ; voir §IV.2 pour la formulation précise), dont la violation paraît évidente (voir l'article IX pour une discussion plus approfondie). Dans le cadre formel d'Anscombe-Aumann, c'est – toujours selon la lecture habituelle – l'axiome d'indépendance qui joue ce rôle (voir §I.3.2 ci-dessous ou §III.2.1 pour la formulation précise). Ainsi, la littérature sur *l'ambiguïté*, qui est large-

4. Ici j'utilise le terme dans un sens assez vague, et d'ailleurs plus conforme à Knight, et non pour me référer à un modèle précis, comme celui de Bewley (1989) qui est aujourd'hui parfois associé au nom de Knight.

ment motivée par les exemples ellsbergiens, peut être organisée selon l'affaiblissent retenu pour l'un des deux axiomes. Le principe de la chose sûre et l'indépendance garantissent effectivement l'existence d'une représentation additivement séparable par rapport aux états (voir §I.2 ci-dessous et l'article IV), et comme c'est la propriété clef mise en cause par les exemples d'Ellsberg, il est naturel qu'elle soit au centre des préoccupations.

Cette approche comportementale repose cependant sur l'hypothèse que, pour comprendre les choix des agents, il faut raisonner sur des relations de préférences complètes – c'est-à-dire telles qu'entre tout couple d'options, le décideur préfère (faiblement) l'une à l'autre. En présence de préférences complètes (et transitives), affaiblir l'axiome d'indépendance est la seule manière de rendre compte des comportements ellsbergiens. Mais il est ici possible de préserver l'indépendance en sacrifiant la complétude. Ainsi, on a proposé des théories, inspirées par la même motivation de rendre compte des situations où le décideur a du mal à former des croyances probabilistes précises, dans lesquelles on affaiblit l'axiome de complétude mais garde l'indépendance.

Alors qu'elles sont parfois présentées comme concurrentes (voir notamment Levi (1986)), ces deux approches sont largement complémentaires. Sur le plan technique, on le sait depuis plus d'une décennie, les travaux de Nehring (2001) et Ghirardato et al. (2004) ayant ouvert la voie. D'un côté, on peut retrouver des préférences incomplètes satisfaisant l'indépendance au sein des relations de préférence ambiguë, c'est-à-dire en tant que sous-relations ; de l'autre, ces dernières peuvent être retrouvées à partir des premières grâce à des procédures de complétion appropriées. Sur le plan conceptuel, Gilboa et al. (2010) ont thématisé le rapport en termes d'une distinction entre rationalité objective et subjective ; pour une discussion d'autres approches, voir §VII.1.3.

Le défi Si l'espérance d'utilité ne convient pas en cas d'incertitude importante, vue la violation de l'indépendance ou de la complétude, alors quelle théorie de décision devrait-on adopter ? À l'heure actuelle, les propositions ne manquent pas. Depuis les travaux novateurs de Schmeidler (1989) et Gilboa and Schmeidler (1989), qui ont introduit et axiomatisé les modèles d'espérance d'utilité à la Choquet et de maximin des espérances d'utilité, plusieurs modèles de choix dans l'ambiguité ont été développés ; pour ne prendre que quelques exemples, Ghirardato et al. (2004); Klibanoff et al. (2005); Maccheroni et al. (2006); Cerreia-Vioglio et al. (2011); Gajdos et al. (2008); Chateauneuf and Faro (2009); Chateauneuf et al. (2007); Tversky and Kahneman (1992); Wakker (2010); Hansen and

Sargent (2001). Quant aux théories qui affaiblissent la complétude, la variété est moins grande ; on commence à développer d'autres modèles que celui de Bewley (2002), qui est encore la référence (voir notamment la discussion dans §VII.1.3). Fidèle à l'exemple donné par le théorème de Savage, on cherche généralement, pour chaque théorie proposée, une caractérisation axiomatique avec une identification (c'est-à-dire une unicité appropriée) de ses composantes ; pour la grande majorité de celles suscitées, et pour d'autres, on a fourni les résultats désirés. Cette richesse remarquable, qui a permis une compréhension raffinée de l'ensemble des possibilités théoriques, nous renvoie à la question initiale : quelle théorie de décision adopter pour les situations d'incertitude knightienne ?

Une question, aussi naïve qu'elle soit, n'est jamais insatisfaisante en soi ; ce ne sont que des présupposés sur les réponses qui nous égarent. Il convient donc de remarquer que rien ne garantit que la nôtre puisse recevoir une réponse univoque. La théorie de l'espérance d'utilité était censée remplir trois rôles, et l'a fait avec plus ou moins de succès. Dès qu'on accepte de s'en éloigner, il faut également admettre la possibilité qu'une seule et même théorie ne soit pas capable de bien remplir les trois rôles. Il se peut qu'une théorie qui décrit bien le comportement des agents dans des situations d'incertitude ne soit pas adéquate normativement, et, par ailleurs, ne soit pas suffisamment simple ou commode au point de vue de la modélisation économique. D'ailleurs, un bref parcours des littératures touchant à l'ambiguïté et consacrées aux dimensions descriptive et de la modélisation donne des raisons de douter de l'existence d'une théorie capable de jouer ces deux rôles, sans même parler du troisième. En outre, il se pourrait qu'il n'existe même pas de théorie qui remplisse un seul des rôles correctement : différents individus dans des situations distinctes pourraient être représentés par des modèles divers, ou certains modèles pourraient être plus commodes dans un type d'application et d'autres plus faciles à utiliser ailleurs.

Pour avancer, il faut donc raffiner la question. Dans les travaux qui suivent, l'accent est mis sur le rôle le moins bien représenté dans la littérature actuelle sur la décision dans l'incertain. Comparée aux nombreux travaux descriptifs, et à la littérature importante et croissante sur les applications, la question normative est peu traitée explicitement en ce moment. Pourtant, elle n'est pas de moindre intérêt. Si les décideurs d'aujourd'hui et de demain doivent prendre des décisions en situations d'incertitude extrême, et si l'économiste veut informer, sinon aider à prendre, ces décisions, il se doit de commencer par une théorie normative ou au moins un approfondissement normatif qui peut servir de guide. Alors que les travaux présentés pourraient être pertinents sur le plan descriptif ou pour des applications, la question qui les motive est normative : quelle théorie – ou théories – de la décision devrait-on adopter face à l'incertitude knightienne ?

I.2 La dépendance de l'utilité à l'égard des états

Pourrait-on trouver un théorème de représentation qui serve de fondement à l'utilité dépendante des états, comme les théorèmes classiques de Ramsey, de Savage et d'Anscombe-Aumann l'ont fait pour l'utilité indépendante des états ? Il est connu de longue date qu'enlever les axiomes pertinentes ne fait pas l'affaire. Par exemple, chez Anscombe-Aumann, si l'on supprime l'axiome de monotonie en gardant les autres, on obtient la représentation suivante :

$$V(f) = \sum_{s \in S} U(s, f(s)) \tag{I.3}$$

où la fonction U, qu'on peut appeler *évaluation*, a des propriétés d'unicité raisonnable (voir par exemple Karni et al. (1983) et l'article II). Pourtant, il y a beaucoup de manières de "factoriser" l'évaluation U en une utilité dépendant des états et une mesure de probabilité : la question est de savoir laquelle est la bonne. En réponse à cette question, il faut un théorème qui fournisse une représentation plus fine que (I.3), séparant U en une probabilité et une utilité qui soient uniques en un sens approprié. Comme l'a souligné Karni (1996, 2011), ce défi se présente même quand les axiomes habituels d'indépendance de l'utilité à l'égard des états sont satisfaits. Sous ces conditions, les théorèmes de représentation classiques assurent qu'il existe une représentation par une fonction d'utilité indépendante des états et une mesure de probabilité et que ces deux sont uniques par rapport à la propriété annoncée de la fonction d'utilité. Ils n'éliminent pas la possibilité qu'il y ait d'autres représentations dans lesquelles la fonction d'utilité soit dépendante des états. Si, avec Karni (voir les articles cités et les références qu'ils donnent), on considère la représentation standard de l'espérance d'utilité comme une famille parmi toutes les représentations comportant une probabilité et une utilité qui puisse dépendre des états, ni l'utilité ni la probabilité ne sont en fait uniques.

Il est commode de poser le problème précédent comme suit : quels axiomes doit-on ajouter à ceux qui donnent la représentation (I.3) pour fournir la séparation désirée de U ? *Grosso modo*, deux genres d'approche sont explorés dans la littérature.

I.2.1 Les actes "subjectivement" constants

Si les axiomes d'indépendance de l'utilité à l'égard des états – disons l'axiome de monotonie chez Anscombe-Aumann – sont violés, alors les actes constants ne sont pas perçus comme tels par le décideur. Recevoir €100 quel que soit l'état de la nature, bien que la conséquence soit la même dans tous les états de la nature, n'est pas un acte "subjectivement" constant si le décideur n'associe pas la même utilité aux €100 dans un état (il est vivant) et dans un autre (il meurt). Alors on peut se demander s'il y a des actes qui, sans nécessairement donner les mêmes conséquences dans chaque état, sont tels que l'utilité attribuée par le décideur aux conséquences est constante dans les différents états. Ainsi, l'acte qui donne €200 si l'agent est vivant et €100 sinon conviendrait si l'utilité de €200 quand il est vivant était égale à celle de €100 en cas de mort. Karni (1993b,a) a montré que, en présence d'un ensemble suffisamment riche de tels actes "subjectivement" constants – qu'il appelle des actes à valuation constante –, on peut formuler un théorème de représentation pour l'utilité dépendant des états, en remplaçant les actes constants ordinaires dans les théorèmes classiques par ces nouveaux actes. Il reste à savoir quand un tel ensemble d'actes existe.

Dans "When is there state independence ?" (l'article II) je propose une réponse, sous la forme d'une condition, exprimée sur les comportements, nécessaire et suffisante pour l'existence de l'ensemble d'actes recherché. À partir de cette caractérisation, en ajoutant des axiomes d'indépendance cardinale au sens indiqué dans §I.1.3, mais formulés en termes d'actes "subjectivement" constants, on peut facilement produire des théorèmes de représentation pour l'utilité dépendant des états qui aient les bonnes propriétés. Le résultat vaut pour les deux cadres formels principaux, mais il s'applique seulement quand l'espace des conséquences (ou des résultats dans le cadre formel d'Anscombe-Aumann) est fini ou dénombrable.

Dans la mesure où la stratégie adoptée consiste à appliquer les théorèmes existants sur l'espérance d'utilité, mais en utilisant un autre ensemble d'actes à la place des actes constants, on peut lui opposer la même objection qu'à ces théorèmes : à savoir, que l'utilité et la probabilité sont uniques seulement sous l'hypothèse que les actes utilisés à la place des actes constants ont vraiment une utilité constante à travers les états. Sans cette hypothèse, on peut répéter l'argument de Karni et obtenir une famille entière d'utilités et de probabilités qui représentent les mêmes préférences. L'article II propose, sous certaines hypothèses, un résultat d'unicité pour l'ensemble d'actes "subjectivement" constants (Corollary II.1), qui peut servir de réponse à cette critique : parmi les multiples représentations, une se distingue en étant la seule pour laquelle il existe des actes "subjectivement" constants. En outre, puisque la condition de l'existence d'un ensemble d'actes "subjectivement" constants est comportementale, dès lors que les préférences la satisfont, on sait que toute représentation par une fonction d'utilité dépendant des états et une autre mesure de probabilité peut être transformée en la représentation de l'article (voir Karni (1996, 2011) pour une transformation similaire). Considérons manitenant l'autre stratégie pour relever le défi posé par la dépendance de l'utilité à l'égard des états.

I.2.2 L'ajout de "données supplémentaires"

Si les préférences ordinaires sur les actes ne suffissent pas à fournir la probabilité et l'utilité au niveau d'unicité recherchée, peut-être faut-il augmenter les données initiales de choix, de situations de choix ou de préférence. Cette stratégie a une longue histoire et peut prendre plusieurs formes : pour ne prendre que quelques exemples, Luce et al. (1971) et Fishburn (1973) prennent les préférences conditionnelles comme primitives et obtiennent une représentation qui ressemble à une dépendance d'utilité à l'égard des états, Karni et al. (1983) se servent de préférences hypothétiques (voir aussi Karni and Mongin (2000)), et Karni (2011) introduit un cadre formel comportant des actions, des effets, des paris et des observations, et considère le rapport entre les préférences avant et après l'observation.

La contribution de "Living without state-independence of utilities" (l'article III) s'inscrit dans cette tradition. Partant de la réponse de Savage dans sa correspondance avec Aumann, l'article développe l'idée que, alors qu'il y a sans doute des situations où l'utilité dépend des états, il y a également des situations – avec d'autres états par exemple – où elle est indépendante. L'utilité de €100 peut varier selon l'état de santé du décideur – donc dans une situation où les états de la nature pertinents pour la décision, étant donné le niveau de finesse nécessaire pour la formaliser, comprennent l'état de santé, l'utilité dépend de l'état. Par contre, l'utilité de €100 ne dépend pas du cheval qui gagne le 15h15 à Longchamp ce samedi – et donc dans le contexte d'une décision sur le pari à faire sur cette course, où les états pertinents spécifient son résultat, l'utilité est indépendante par rapport à ces états. Mais s'il y a de telles situations, où l'utilité ne dépend pas des états (et donc où les axiomes tels la monotonie sont satisfaits), pourquoi pas s'en servir pour révéler les probabilités et les utilités du décideur ? Dans le cadre formel d'Anscombe-Aumann, cet article identifie des conditions relatives à des situations différentes qui suffisent pour révéler une fonction d'utilité unique (à une transformation affine près) et une mesure de probabilité unique qui permettent de représenter les préférences du décideur. Il introduit un élément nouveau dans la représentation – un facteur dépendant de la situation (*situation-dependent factor*) – qui traduit les phénomènes comportementaux correspondant à la dépendance de l'utilité par rapport aux états. Ce facteur permet une distinction entre l'utilité absolue, qui est indépendante des états (et donc peut être employée dans n'importe quelle situation de décision, quel que soit l'espace d'états pertinent) et l'utilité "relative à la situation", issue de la combinaison de l'utilité absolue et du facteur en question. (Voir §III.3 pour une discussion du rapport avec l'utilité dépendante des états habituelle ainsi que des autres approches proposées dans la littérature.)

La dernière section de l'article (§III.4) montre comment éviter la notion de situation, en formulant un résultat semblable dans une seule "grande" situation, où l'indépendance par rapport aux états est cherchée dans des "fragments" (des sous-algèbres de l'ensemble d'événements, par exemple).

I.2.3 La dépendance de l'utilité à l'égard des états dans le cadre de Savage

Les articles précédents partent de la représentation (I.3) : le premier la suppose d'emblée, sans poser la question de sa caractérisation comportementale ; le second travaille dans le cadre formel d'Anscombe-Aumann, où une telle caractérisation est connue. Toute extension de ces résultats, et d'ailleurs de certains autres dans la littérature (par exemple, Karni et al. (1983)), au cadre formel de Savage nécessitera de formuler un équivalent de (I.3) pour un ensemble de conséquences dénué de structure et pour un espace d'états infini, avec la caractérisation axiomatique afférente. Le dernier article portant sur la dépendance de l'utilité à l'égard des états, "An additively separable representation in the Savage framework" (l'article IV), traite de cette question. Le théorème qu'il propose obtient la représentation suivante, qui peut être considérée comme l'équivalent de (I.3) dans le cadre de Savage (voir

§IV.3) :

$$V(f) = \int_f dU \tag{I.4}$$

où U est une mesure (finiment additive) sur l'espace produit $S \times C$ (avec une σ-algèbre appropriée). Le résultat est largement sans surprise : les axiomes comportementaux sont ceux de Savage (moins les axiomes d'indépendance de l'utilité) et les axiomes techniques sont des versions des axiomes de continuité proposés ailleurs dans la littérature, plus un axiome spécifique. Comme discuté dans l'article, ce résultat peut s'appliquer au-delà de la théorie de la décision dans l'incertitude, lorsqu'on veut une représentation additivement séparable avec un ensemble d'indices continu (représentant le temps, par exemple) mais sans hypothèses sur la structure des conséquences.

I.3 L'incertitude knightienne

Comment devrait-on choisir dans des situations d'incertitude extrême, notamment celles où il est difficile de former des croyances probabilistes ? Parmi les questions qui constituent le défi de l'incertitude knightienne, celle-ci a dirigé mes recherches sur le sujet. Bien que la plupart des travaux présentés se concentrent sur l'incertitude relative à l'état de la nature – le côté de la croyance visé par la distinction knightienne –, il convient de remarquer que "l'incertitude" d'un décideur peut également planer sur ses propres préférences ou utilités (si l'on peut toujours utiliser ces termes dans les cas où elles ne sont éventuellement même pas formées).

Pour certains, la question qui motive ces travaux est superflue, car on connaît de longue date la réponse : quels que soient ses défauts en tant que description du comportement des individus, la théorie de l'espérance d'utilité fournit toujours la norme que les décideurs devraient viser. Certains économistes (Gilboa et al., 2009, 2012), ainsi que plusieurs philosophes (Levi, 1986; Bradley, 2009; Joyce, 2011), ont exprimés leur scepticisme à cet égard. Le projet guidant l'essentiel de cette partie est de développer et défendre une théorie de la décision qui pourrait concurrencer l'étalon bayésien sur le plan normatif.

Une pareille théorie doit alors rend compte de comportements, dont les fameux exemples ellsbergiens mentionnés ci-dessus, qui peuvent être considérés comme acceptables du point de vue normatif. La proposition élaborée et explorée dans les travaux présentés est le fruit de l'intuition que la confiance que nous avons dans nos attitudes – qu'elles soient des désirs (ou utilités) ou des croyances (ou jugements de probabilité) – mérite d'avoir un rôle dans la décision. La maxime inspiratrice est la suivante : plus importante est la décision, plus on exige de confiance dans les jugements sur lesquels on s'appuie pour la prendre.

I.3.1 Le cadre général et premières contributions

"Confidence and Decision" (l'article V) développe le cadre général, discute certaines de ses propriétés, et en explore les conséquences pour le choix dans l'incertain. Comme il a été remarqué ci-dessus, la littérature sur l'ambiguïté se concentre sur des violations de l'axiom d'indépendance (dans le cadre d'Anscombe-Aumann) et suppose généralement une seule fonction d'utilité indépendante des états. Il va de même dans cet article, qui concentre entièrement sur le rôle de la confiance dans les croyances, et propose un modèle qui viole l'axiome d'indépendance. Celui-ci introduit trois nouveaux éléments :

— une représentation des croyances et de la confiance dans les croyances par un *confidence ranking* ("classement par confiance") [5] – une famille emboîtée d'ensembles de mesures de probabilité – qu'on dénote par Ξ. Chaque ensemble dans la famille correspond à un niveau de confiance et représente les croyances dans lesquelles le décideur a, au moins, autant de confiance (voir §§V.1.3 et V.2.1).

— le *cautiousness coefficient* ("coefficient de sagesse") – une allocation d'un ensemble dans le *confidence ranking* à chaque décision – qu'on dénote par D. Cette fonction donne le niveau de confiance approprié pour une décision, qui dépend de l'importance de la décision ou de ses enjeux. Cette composante endogène du modèle représente le goût du décideur pour prendre des décisions sur la base d'une confiance limitée.

— les enjeux dans une décision, ou en d'autres termes son importance. Cet article suppose une notion particulière d'enjeu, qui est la pire conséquence possible si l'on choisit l'acte. À la différence des éléments précédents, celle-ci est exogène au modèle, et constitue un facteur sur lequel différents modèles divergent.

5. Pour faciliter la mise en rapport de cette note et les articles qui sont rédigés en anglais, je continue d'utiliser les terme anglais par la suite.

Dans la représentation axiomatisée, les préférences sont représentées par la fonctionnelle suivante sur les actes, où σ est la notion d'enjeu donnée dans §V.2.1 :

$$V(f) = \min_{p \in D(\sigma(f))} \sum_{s \in S} p(s)u(f(s)) \tag{I.5}$$

Ce modèle est une extension du modèle maximin EU de Gilboa and Schmeidler (1989), la différence axiomatique entre les deux étant par ailleurs assez faible. En gros, l'affaiblissement de l'indépendance respecté par le modèle maximin EU peut s'interpréter comme exigeant que la préférence entre deux actes ne soit pas sensible aux enjeux impliqués dans le choix entre eux. Bien entendu, du point de vue de l'intuition de départ, cet axiome est trop fort. Le modèle proposé l'affaiblit de la manière qu'on attendrait : si les enjeux n'augmentent pas, alors toute préférence pour un acte "ambigu'" par rapport à un acte qui ne l'est pas est préservée, et de même pour la direction inverse. La relative plausibilité de ces conséquences comportementales – qui, comme mentionné dans §§I.1.1 et I.1.2, constituent un contrôle possible du modèle – pourrait être considérée favorablement dans l'évaluation normative du modèle.

Une propriété bien prisée de la théorie de l'espérance d'utilité est la séparation apparente entre les croyances, représentées par la mesure de probabilité, et les goûts ou les désirs, représentés par la fonction d'utilité. Une séparation aussi nette est beaucoup moins répandue parmi les modèles d'ambiguïté, largement en raison de la difficulté à séparer l'incertitude et l'attitude de choix en présence de cette incertitude. D'ailleurs, le modèle "lisse" (*smooth*) de Klibanoff et al. (2005) est presque le seul qui prétende à cette séparation (voir aussi Epstein (2010); Klibanoff et al. (2012)). Comme suggéré ci-dessus, le modèle ici proposé paraît admettre une telle séparation, dans la mesure où elle comporte une composante représentant les croyances (le *confidence ranking*) et une autre qui reflète une attitude vis-à-vis de l'incertitude (le *cautiousness coefficient*). L'article propose des exercices de statique comparative qui confortent ces interprétations (§V.3). Si ils sont jugés acceptables, ce sera un autre point fort de l'approche, sur le plan normatif.

Comme discuté dans la dernière section de l'article, le cadre proposé permet de définir d'une famille de modèles, dont tous respectent plus ou moins la maxime "plus les enjeux sont élevés, plus il faut de confiance dans les croyances pour se servir d'elles". D'une part, la règle de décision (maximin EU dans (I.5)) peut varier entre les différents membres de la famille ; d'autre part, la notion d'enjeu (qui est respectée par D) peut différer entre des modèles.

Dans "Confidence in Preferences" (l'article VI), l'approche est étendue des croyances aux préférences. L'article est ainsi le seul qui ne traite pas d'incertitude. Il montre comment utiliser la même sorte de représentation (une famille emboîtée d'ensembles, mais de préférences plutôt que de probabilités, un *cautiousness coefficient* et une notion d'enjeu, appelée importance dans ce contexte) pour construire une théorie du choix dans le certain. La maxime de base reste identique, même si sa réalisation est différente : là où le décideur manifestait plus d'aversion à l'incertitude en absence d'une confiance suffisante, ici il peine à former des préférences. D'ailleurs, l'introduction et le traitement de l'incomplétude des préférences est une autre contribution par rapport à l'article précédent. Travaillant dans un cadre où les éléments primitifs sont des fonctions de choix, qui donnent les options choisies dans des menus, il propose une retraduction de l'incomplétude des préférences en termes comportementaux par le fait d'éviter la décision, soit en la transmettant à quelqu'un d'autre soit en l'ajournant. Il sera commode par la suite de se servir de l'anglais *defer*, qui désigne à la fois le fait de déférer une décision à autrui et celui de la différer à un autre moment, et qui reflète mieux la communauté entre ces cas. L'article représente le *deferral* en utilisant des fonctions de choix qui peuvent ne donner aucun élément du menu, et en introduisant un nouvel axiome qui peut être vu comme l'équivalent de l'axiome β de Sen pour les préférences incomplètes. Enfin, l'article approfondit certaines notions et hypothèses du modèle, en particulier celles d'enjeu et de confiance (§§VI.3.1 et VI.3.3).

I.3.2 Développements et défense

I.3.2.1 Questions et défis

Les articles précédents ouvrent plusieurs questions. Si un manque de confiance peut entraîner l'aversion vis-à-vis de l'incertitude, mais aussi l'incomplétude ou le *deferral*, quel est le rapport entre ces deux conséquences ? Plus généralement, si la confiance a vraiment un rôle à jouer dans la décision, on s'attendrait à ce qu'il y ait des conséquences comportementales spécifiques : peut-on en trouver ? Mais plus importants encore que ces questions sont les défis restants. Ils sont des deux types expliqués dans §I.1.1. D'une part, il y un défi "fondationnel". Malgré les arguments proposés dans l'article VI, la notion d'enjeu est plus naturellement interprétée comme une attitude ou une représentation du décideur et, en tant que telle, ne devrait pas être supposée de manière exogène. Elle a donc besoin d'un fondement : un théorème de représentation dans lequel elle est révélée de manière endogène. D'autre part, il y a un défi "d'évaluation". Des arguments connus cherchent à mettre dans l'embarras tout modèle qui ne satisfait pas l'axiome d'indépendance standard en examinant ses conséquences dynamiques. Descriptivement correct ou non, approprié à la modélisation économique ou non, tout modèle d'ambiguïté, y compris celui proposé dans l'article V, serait condamné par ces arguments sur le plan normatif. A ce point qu'ils constituent en réalité la pierre angulaire de la défense de l'espérance d'utilité comme théorie normative de la décision dans l'incertain. Tout projet de développer une autre théorie reste incomplet si l'on ne les reconsidère pas.

I.3.2.2 Préférences incomplètes, choix et deferral

Les articles VII and VIII cherchent à répondre aux questions ouvertes et au premier défi mentionnés ci-dessus. "Incomplete Preferences and Confidence" (l'article VII) explore les effets de la confiance sur l'incomplétude des préférences en la rapportant au choix et à l'ambiguité, cela dans un cadre standard de décision dans l'incertain. La représentation obtenue est de la même forme générale que (I.5), mais utilise la règle d'unanimité de Bewley plutôt que le maximin EU : pour tout couple d'actes f et g, $g \succeq f$ si et seulement si :

$$\sum_{s\in S} p(s)u(g(s)) \geqslant \sum_{s\in S} p(s)u(f(s)) \qquad \forall p \in D((f,g)) \tag{I.6}$$

Ainsi, la représentation appartient à la famille proposée dans l'article V. Du point de vue comportemental, les axiomes sont tout à fait raisonnables dès qu'on introduit une nouvelle distinction entre la *violation* et l'*abstention* d'un axiome. Pour illustrer, prenons le cas de l'axiome d'indépendance, qui demande que pour tout couple d'actes f et g et toutes mixtures $f_\alpha h$, $g_\alpha h$ avec un autre acte h, $f \preceq g$ si et seulement si $f_\alpha h \preceq g_\alpha h$. (Voir §VII.2 pour des détails.) Ainsi l'axiome exige que si un décideur préfère (faiblement) g à f, il ne peut pas préférer (fortement) $f_\alpha h$ à $g_\alpha h$, mais il ne lui est pas permis non plus d'avoir des préférences indéterminées entre ces deux actes. La version de l'indépendance proposée dans l'article VII – appelée *l'indépendance pure* – rejoint l'axiome classique sur le premier point, mais diverge sur le deuxième. Elle condamne toute "violation" de l'axiome de l'indépendance : c'est-à-dire, toute combinaison de préférences déterminées qui sont incompatibles avec l'axiome. Par contre, il permet des "abstentions" de l'axiome : des combinaisons de préférences qui sont indéterminées sur des couples d'actes où l'axiome standard interdit l'incomplétude.[6] Sur le plan de la plausibilité, l'indépendance pure n'a rien à envier à l'axiome classique : les combinaisons de préférences (généralement déterminées) qui sont récusées par l'axiome standard le sont aussi par cette nouvelle version.[7] Au contraire, l'article propose des exemples où les deux axiomes divergent, et où seul le traditionnel est critiquable.

L'article touche aussi au rapport entre préférences incomplètes et choix, et il tire certaines conséquences du modèle pour les marchés financiers. (Remarquons au passage qu'il reste relativement silencieux sur l'interprétation de l'incomplétude, et que les résultats peuvent être lus dans des optiques descriptive et de modélisation économique, aussi bien que normative.) Tout d'abord, la mise en jeu de la confiance met en lumière deux stratégies générales pour décider quand les préférences sont incomplètes (§VII.4). L'une retient uniquement l'usage des croyances dans lesquelles le décideur a suffisamment confiance, en utilisant une règle autre que celle impliquée dans (I.6) pour décider. L'autre permet au décideur de mobiliser toutes ses croyances, y compris celles dans lesquelles il a peu de confiance. À ma connaissance, cette distinction n'a pas été introduite clairement auparavant dans la littérature. C'est un avantage de la notion de confiance qu'elle la fait bien ressortir. Pour illustrer encore l'intérêt de cette notion, l'article aborde ses conséquences dans des marchés financiers (§VII.5). Comparé au modèle standard de Bewley (2002), le modèle intégrant la confiance ne diffère pas tant dans la forme des optima de Pareto que dans la difficulté à les atteindre. La présence de la confiance – qui peut empêcher des décideurs d'accepter des gros échanges alors qu'ils acceptent des petits – peut conduire à des allocations qui ne sont pas des optima de Pareto mais où l'on ne peut pas facilement accéder à ceux qui les auraient dominés. Cette inertie ou ce conservatisme dus à la confiance n'apparaît pas dans les modèles courants (bayésiens ou non) des marchés financiers.

L'article VIII, "Confidence as a Source of Deferral", peut être considéré comme un

6. Il est évident que cette distinction s'écroule quand les préférences sont complètes.

7. En particulier, les arguments dynamiques qui critiquent les violations de l'indépendance ne touchent pas à ce modèle, ce qui montre que les problèmes normatifs relevant de ces arguments ne sont pas inhérents à la famille de modèles proposée dans l'article V, mais au pire concernent certains de ses membres. Comme il sera clair par la suite, alors que c'est une réponse au défi posé par ces arguments, ce n'est pas celle que je prônerai.

développement technique du précédent qui explore quelques notions mises en avant dans l'article VI, notamment celle de *deferral*. Sur le plan conceptuel, sa contribution est d'introduire un autre rôle pour la confiance que celui développé antérieurement. Non seulement on peut vouloir éviter de décider en absence de confiance, par le *deferral* par exemple, mais le manque de confiance peut être relié qualitativement à la valeur que le décideur attribue à l'option de ne pas devoir décider. Cette idée peut avoir des conséquences économiques significatives. La grande majorité de la littérature sur le risque et l'information, en théorie des contrats par exemple, néglige le rôle de la confiance comme raison éventuelle pour différer un choix ou le déférer à autrui, en modélisant ces décisions uniquement par la différence d'information (concernant le monde, mais peut-être aussi les préférences futures) entre agents. Les modèles existants de préférences incomplètes ont du mal à prendre en compte le coût du *deferral* dans ces situations : alors qu'ils peuvent déterminer *quand* le décideur ne veut pas décider, mais ils ne peuvent pas dire *combien* il est prêt à payer pour ne pas devoir décider. L'article VIII cherche à combler ce manque en développant un modèle qui intègre la confiance mais qui s'applique dans ces situations.

Outre sa contribution sur les plans conceptuel et économique – il développe non seulement un modèle de *deferral*, mais esquisse une manière d'intégrer la confiance dans l'approche classique basée sur la différence d'information – l'article apporte une solution au problème de fondements mentionné ci-dessus. Aucune notion d'enjeu ou d'importance n'est maintenant supposée de manière exogène, et alors que le modèle n'explicite pas ces notions, il est possible de le reformuler à l'aide d'une notion d'enjeu, qui serait alors révélée de manière endogène (voir en particulier Remarques VIII.1, VIII.3 et §VIII.5). On peut ainsi considérer que le théorème de représentation proposé ici fonde cette notion d'enjeux, ce qui manquait aux articles précédents.

I.3.2.3 Le choix dynamique

Le dernier défi est celui posé par les arguments qui prétendent démontrer que tout modèle violant l'axiome d'indépendance (ou le principe de la chose sûre) est voué à avoir des conséquences indésirables dans des situations dynamiques : il découlerait une violation de l'un ou l'autre des principes dynamiques raisonnables, par exemple, le décideur refuserait de l'information gratuite ou même choisirait des options dominées. Ces arguments mettent donc en doute les prétentions normatives des modèles d'ambiguïté. La famille de modèles esquissée dans l'article V n'est pas touchée dans sa totalité par ce genre d'argument, du moins dans sa forme la plus grave, car certains membres, notamment ceux avec des préférences incomplètes, ne violent pas l'axiome d'indépendance (voir la discussion de l'article VII ci-dessus).

Pourtant, ce ne serait pas une manière adéquate de défendre la validité normative de l'approche. Si l'ambition était de proposer une manière "rationnelle" de décider dans des situations d'incertitude extrême, où les jugements exacts de probabilité sont difficiles à former mais où il faut néanmoins faire un choix, une théorie qui ne s'applique pas dans ces situations précisément ne serait guère utile. "Dynamic Consistency and Ambiguity : A Reappraisal" (l'article IX) cherche donc à réfuter ces arguments. Dans la mesure où ils touchent à tout modèle violant l'axiome d'indépendance, la réfutation, si elle réussit, défendra non seulement l'approche élaborée ici mais toute théorie de l'ambiguité.

Un des arguments clefs consiste en un résultat mathématique portant sur deux principes dynamiques : le *conséquentialisme*, qui demande que les préférences du décideur à un moment et dans une situation donnée soient indépendantes du passé et de ce qui aurait pu arriver, et la *cohérence dynamique*, qui exige que ses préférences *ex ante* concernant les plans pour les éventualités futures (*contingent plans*) soient en harmonie avec ses préférences *ex post*. Le résultat en question affirme que, sous l'hypothèse de quelques propriétés de base, toute violation de l'axiome d'indépendance entraîne une violation de l'un ou l'autre de ces principes, ce qui, vu leur plausibilité, poserait un problème sur le plan normatif (voir §IX.1). Certains autres problèmes soulevés à propos des modèles d'ambiguïté sont liés à cet argument (voir par exemple la discussion de l'aversion à l'information, §IX.5).

L'article IX part du constat banal que tout principe dynamique concernant les éventualités, telle la cohérence dynamique, devrait être formulé par rapport aux éventualités que le décideur lui-même envisage, qui sont, après tout, les seules qu'il prenne en compte dans la formation de ses plans. Pourtant, dans la formalisation traditionnelle de la cohérence dynamique, il n'y a aucune considération des éventualités envisagées par le décideur. Par conséquent, le résultat du paragraphe précédent s'appuie, au pire, sur une formalisation erronée du principe de cohérence dynamique (et donc ne pose aucun problème sur le plan normatif) ; au mieux, il repose sur une hypothèse cachée concernant les éventualités envisagées (à savoir, qu'elles coincident avec celles posées par le théoricien). L'article décèle les implications de cette hypothèse pour montrer que, dès qu'elle est prise en compte, les prétendus embarras pour les théories qui violent l'indépendance s'avèrent illusoires. D'un côté, en l'absence de l'hypothèse, l'incompatibilité entre les principes dynamiques et la violation de l'indépendance disparaît (§IX.3). De l'autre côté, sa présence implique une forme spéciale des croyances *ex ante*, et il suffit de respecter cette implication pour dissoudre les prétendus problèmes pour les modèles d'ambiguïté. D'ailleurs, le constat qu'il y a des implications propres à l'hypothèse fournit et justifie un conseil concret pour l'usage de ces modèles dans des situations de choix dynamique (§IX.4). Enfin, l'article étaie la notion d'éventualité envisagée par le décideur, notamment sur le plan "fondationnel", avec un théorème de représentation (§IX.6).

I.4 Au croisement de l'indépendance de l'utilité à l'égard des états et de l'incertitude knightienne

Chacun des travaux précédents, comme la littérature à laquelle ils appartiennent, concentre sur l'une des deux objections à la théorie de l'espérance de l'utilité discutées précédemment, en faisant abstraction de l'autre. Or il y a manifestement des situations impliquant à la fois la dépendance de l'utilité par rapport aux états et l'incertitude knightienne – l'exemple déjà évoqué du choix de politique environnementale en est probablement une. Le dernier article présenté, "Uncertainty aversion, multi utility representations and state independence of utility" (l'article X) tente de prendre en compte les deux objections, en développant un modèle de décision qui ne suppose ni la monotonie ni l'indépendance, mais qui représente un décideur manifestant de l'aversion par rapport à l'incertitude. Il montre que cette propriété bien connue dans la littérature sur l'ambiguïté, en présence de certains axiomes de base et quelques conditions techniques, caractérise la représentation suivante :

$$V(f) = \min_{U \in \mathcal{U}} \sum_{s \in S} U(s, f(s)) \tag{I.7}$$

où $\mathcal{U}$ est un ensemble d'évaluations. Cette représentation généralise plusieurs modèles dans la littérature sur l'ambiguïté, dont celui de Gilboa and Schmeidler (1989) qu'elle évoque directement, mais aussi la représentation (I.3), qui est un point de départ pour la littérature sur l'utilité dépendante des états (voir §I.2). Par ailleurs, elle aussi peut servir comme base, mais cette fois pour le développement d'une théorie de l'utilité dépendante des états en présence de l'aversion par rapport à l'incertitude : pareillement au cas considéré dans §I.2, la tâche serait de trouver une séparation convenable de l'ensemble $\mathcal{U}$ en un facteur d'utilité et en un autre qui a trait aux probabilités. Remarquons qu'une telle séparation produirait éventuellement un ensemble de fonctions d'utilité plutôt qu'une seule fonction. De tels ensembles peuvent s'interpréter par l'*imprécision* des désirs du décideur, de la même manière que les ensembles de probabilités chez Gilboa and Schmeidler (1989) et Bewley (2002) représenteraient des croyances qui manquent d'être des probabilités précises. Il est parfois difficile de comparer les conséquences pertinentes – les décisions de politique environnementale, qui demandent souvent de mettre en balance le bien-être de générations différentes, fournissent de nouveau un exemple –, et la représentation (I.7) a le mérite de continuer à s'appliquer dans ces cas.

Il aurait été naturel de développer une théorie de l'utilité dépendante des états dans l'ambiguïté à partir de représentation (I.7), mais l'article X ne le fait pas. Il met plutôt en évidence une équivoque dans la notion d'utilité *indépendante* des états, qui surgit dès que l'on renonce à l'hypothèse d'une seule fonction d'utilité. L'article identifie et caractérise deux notions d'indépendance de l'utilité par rapport aux états : l'une est l'analogue de celle qui est en passe de s'imposer dans la littérature sur les préférences incomplètes (Galaabaatar and Karni, 2013), et l'autre est impliquée par l'axiome de monotonie évoqué ci-dessus. La découverte étonnante est *grosso modo* que ces notions sont incompatibles quand les désirs sont imprécis. Ce qui laisse penser que le croisement entre les questions de l'indépendance de l'utilité par rapport aux états et celles de l'incertitude ou de l'imprécision relative aux croyances ou aux désirs est peut-être plus complexe que l'on n'a imaginé jusqu'à présent.

I.5 Perspectives de recherche

Les travaux présentés dans §I.3 constituent le début d'un programme de recherche plutôt que son aboutissement. Dans les années à venir, j'ai l'intention de développer le programme en plusieurs directions, dont voici certaines.

La première, qui est plutôt de l'ordre du renforcement que de l'extension, concerne les prétentions normatives des modèles proposés. Il s'agira d'examiner de plus près, et avec un œil plus réflexif, ces prétentions, en se confrontant à la littérature philosophique sur la décision et la rationalité, par exemple. Une défense sur ce plan passera non seulement par l'argumentation conceptuelle, mais également par la participation à certains débats actuels pertinents. Le plus évident porte sans doute sur le fameux Principe de Précaution, qui, alors qu'il implique des notions non évoquées ici, comme l'irréversibilité, est le genre de questions sur lesquelles on s'attendrait à une contribution de la part des théoriciens de la décision travaillant dans une perspective normative. Un développement dans cette direction ira vers des préoccupations pratiques, en regardant les procédures utilisées par des scientifiques ayant trait à des incertitudes extrêmes (climatologues, géographes, ingénieurs, statisticiens) pour voir si le cadre développé ici aurait quelque chose à offrir.

Une autre direction consistera à développer une théorie dynamique pour la représentation de confiance proposée. L'article IX ouvre la voie, dans la mesure où il laisse ouverte la possibilité d'une théorie normativement plausible de la dynamique. Une piste possible part de l'idée que la confiance a un rôle non seulement dans la décision mais également dans l'apprentissage ou le changement des croyances. Encore une fois, une maxime peut nous guider. Quand on change de croyances, on retient certaines "veilles" croyances mais souvent on en supprime d'autres pour permettre l'ajout des nouvelles. Il est raisonnable que les croyances qui sont retenues soient celles dans lesquelles on a le plus confiance, et qu'on retire seulement celles dans lesquelles on a le moins confiance. Appuyant sur les développements dans le cadre non-probabiliste de la révision des croyances en logique, philosophie et informatique (voir par exemple Hill (2008) et les références là-dedans), où l'on trouve une idée voisine, le but serait de proposer une théorie du changement de la croyance et de la confiance pour la famille de modèles proposée dans l'article V, de la caractériser du point de vue comportemental et d'en explorer les conséquences.

Une troisième direction, plus expérimentale, consistera à développer des procédures pour révéler les composantes du modèle – en particulier la confiance dans les croyances (le *confidence ranking*) et l'attitude à choisir sur la base d'une confiance limitée (le *cautiousness coefficient*). Une question qui se pose naturellement concernant des éventuelles applications est celle de la *calibration de l'incertitude* (ou des croyances). Les loteries constituent une "échelle" naturelle pour la calibration des croyances probabilistes. En effet, un bayésien peut "situer" chaque croyance par rapport à toute probabilité donnée : par exemple, il peut comparer sa croyance concernant l'éventualité d'une montée de températures globale de plus de 5°C dans les prochains 50 ans à une quantité d'émissions de

CO_2 donnée à la chance que la prochaine boule tirée d'une urne avec 5 boules noires et 95 boules blanches soit noire. De telles comparaisons (ou plutôt leur version "comportementale", en termes de paris) sont par ailleurs importantes pour certaines procédures de révélation des croyances. Pourtant, il n'existe pas d'échelle de ce genre pour les croyances non-probabilistes, et *a fortiori* pour la confiance dans les croyances. L'objectif serait d'en proposer une ; une possibilité serait d'utiliser des urnes telles que le décideur ne connaisse pas la proportion des différentes couleurs mais ait eu la possibilité d'apprendre. (Les attitudes des décideurs dans ce genre de situations sont étudiées expérimentalement dans Abdellaoui et al. (2014).) En comparant les jugements concernant le climat futur à ceux relatifs à la couleur de la prochaine boule tirée dans ces situations – par exemple, que la prochaine boule soit noire après avoir observé trois boules noires et douze boules blanches –, on pourrait calibrer la confiance dans les premiers jugements sur l'échelle formée par les derniers. De plus, en s'appuyant sur l'article VIII, et en particulier sur l'idée que le décideur est moins enclin à différer une décision quand il a plus de confiance, il pourrait même être possible de développer une procédure de choix, dans lequel le prix qu'il est prêt à payer pour différer les décisions révèle sa confiance.

Un certain nombre de ces recherches éventuelles seront entreprises dans le cadre du projet ANR Jeunes Chercheurs "DUSUCA" (Decision Making and Belief Change in the face of Severe Uncertainty : A Confidence-Based Approach), dont je suis le coordinateur, et qui commence en janvier 2015.

Bibliographie

Abdellaoui, M., Hill, B., Kemel, E., and Maafi, H. (2014). Ambiguity Attitudes in Sampled Risk. Mimeo HEC Paris.

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34 :199–205.

Arrow, K. J. (1971). *Essays in the Theory of Risk Bearing*. Markham Publishing., Chicago.

Arrow, K. J. (1974). Optimal Insurance and Generalized Deductibles. *Scandinavian Actuarial Journal*, 1 :1–42.

Bewley, T. F. (1986 / 2002). Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2) :79–110.

Bewley, T. F. (1989). Market Innovation and Entrepeneurship : A Knightian View. Technical Report 905, Cowles Foundation.

Bradley, R. (2009). Revising incomplete attitudes. *Synthese*, 171(2) :235–256.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. (2011). Uncertainty averse preferences. *Journal of Economic Theory*, 146(4) :1275–1330.

Chateauneuf, A., Eichberger, J., and Grant, S. (2007). Choice under uncertainty with the best and worst in mind : Neo-additive capacities. *Journal of Economic Theory*, 137(1) :538–567.

Chateauneuf, A. and Faro, J. H. (2009). Ambiguity through confidence functions. *J. Math. Econ.*, 45 :535–558.

Cozic, M. and Hill, B. (forthcoming). Representation Theorems and the Semantics of Decision-Theoretic Concepts. *Journal of Economic Methodology*.

Drèze, J. H. (1987). *Essays on Economic Decisions under Uncertainty*. Cambridge University Press, Cambridge.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quart. J. Econ.*, 75(4) :643–669.

Epstein, L. G. (2010). A Paradox for the “Smooth Ambiguity” Model of Preference. *Econometrica*, 78(6) :2085–2099.

Fishburn, P. C. (1973). A Mixture-Set Axiomatization of Conditional Subjective Expected Utility. *Econometrica*, 41(1) :1.

Gajdos, T., Hayashi, T., Tallon, J.-M., and Vergnaud, J.-C. (2008). Attitude toward imprecise information. *J. Econ. Theory*, 140(1) :27–65.

Galaabaatar, T. and Karni, E. (2013). Subjective expected utility with incomplete preferences. *Econometrica*, 81(1) :255–284.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2) :133–173.

Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. (2010). Objective and Subjective Rationality in a Multiple Prior Model. *Econometrica*, 78(2) :755–770.

Gilboa, I., Postlewaite, A., and Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms ? *Economics and Philosophy*, 25(3) :285–296.

Gilboa, I., Postlewaite, A., and Schmeidler, D. (2012). Rationality of belief or : why savage's axioms are neither necessary nor sufficient for rationality. *Synthese*.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2) :141–153.

Hansen, L. P. and Sargent, T. J. (2001). Robust Control and Model Uncertainty. *The American Economic Review*, 91(2) :60–66.

Hill, B. (2008). Towards a “Sophisticated” Model of Belief Dynamics. Part II : Belief Revision. *Studia Logica*, 89(3) :291–323.

Joyce, J. M. (2011). A Defense of Imprecise Credences in Inference and Decision Making. In Gendler, T. S. and Hawthorne, J., editors, *Oxford Studies in Epistemology*, volume 4. OUP.

Karni, E. (1993a). A Definition of Subjective Probabilities with State-Dependent Preferences. *Econometrica*, 61 :187–198.

Karni, E. (1993b). Subjective expected utility theory with state dependent preferences. *J. Econ. Theory*, 60 :428–438.

Karni, E. (1996). Probabilities and Beliefs. *J. Risk Uncertainty*, 13 :249–262.

Karni, E. (2011). A theory of Bayesian decision making with action-dependent subjective probabilities. *Economic Theory*, 48(1) :125–146.

Karni, E. and Mongin, P. (2000). On the Determination of Subjective Probability by Choices. *Management Science*, 46 :233–248.

Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51 :1021–1032.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6) :1849–1892.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2012). On the smooth ambiguity model : A reply. *Econometrica*, 80(3) :1303–1321.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston and New York.

Levi, I. (1986). *Hard Choices. Decision making under unresolved conflict*. Cambridge University Press, Cambridge.

Luce, R. D., Krantz, D. H., and Mar, N. (1971). Conditional Expected Utility. *Econometrica*, 39(2) :253–271.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 74(6) :1447–1498.

Nehring, K. (2001). Ambiguity in the context of probabilistic beliefs. Technical report, University of California, Davis.

Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica*, 32 :122–136.

Ramsey, F. P. (1931). Truth and Probability. In *The Foundations of Mathematics and Other Logical Essays*. Harcourt, Brace and Co., New York.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Schmeidler, D. (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3) :571–587.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory : Cumulative representation of uncertainty. *J. Risk Uncertainty*, 5 :297–323.

Wakker, P. P. (2010). *Prospect Theory For Risk and Ambiguity*. Cambridge University Press, Cambridge.

II When is there state independence?

Abstract

Whether a preference relation can be represented using state-independent utilities as opposed to state-dependent utilities may depend on which acts count as constant acts. This observation underlies an extension of Savage's expected utility theory to the state-dependent case that was proposed in this journal by Edi Karni. His result contains a condition requiring the existence of a set of acts which can play the role of constant acts and support a representation involving a state-independent utility function. This paper contains necessary and sufficient conditions on the preference relation for such a set of acts to exist. Results are obtained both for the Savage and the Anscombe and Aumann frameworks. Among the corollaries are representation theorems for state-dependent utilities. Relationships to Karni's work and extensions of the results are discussed. [1]

Keywords: Subjective expected utility; State-dependent utility; Monotonicity axiom

JEL classification: D81, C69.

1. This is slightly modified version of a paper of the same name that appeared in Journal of Economic Theory, 144(3):1119-1134, May 2009.

II.1 Introduction

There has been much discussion of representations of preferences with state-dependent utilities Karni et al. (1983); Karni and Schmeidler (1993); Karni (1993b,a); Karni and Mongin (2000); Drèze (1987); Schervish et al. (1990). In several of these discussions, it has been noted that, by redefining which acts (ie. functions from states to consequences) count as constant, one can transform a state-independent representation into a state-dependent one, and conversely, a state-dependent representation into a state-independent one. To take the example proposed in Schervish et al. (1990), if the acts are bets on the exchange rates between dollars and yen, then a representation which is state-independent when the stakes are formulated in one currency will be state-dependent when the stakes are formulated in the other. Indeed, the idea that failures of Savage's state-independence axioms may come about because the consequences do not yield the acts which are "really" constant is behind some theories of state-dependent utility. Most notably, the extension of Savage's expected utility theory Savage (1954) proposed several years ago by Edi Karni in this journal Karni (1993b), and the extension he proposed of Anscombe and Aumann's version of this theory Anscombe and Aumann (1963); Karni (1993a) rely precisely on this idea. In particular, he introduces the notion of *constant valuation acts*, which, although they are not constant acts, play the role of constant acts in Savage- or Anscombe and Aumann-like representation theorems: state–independence holds with respect to the constant valuation acts instead of the constant acts.

Under what conditions does there exist a set of acts that can play the role of constant acts and that yield a state-independent utility representation of this sort? In Karni (1993b), Karni simply poses the existence of such a set of acts as a condition in his theorem. However, standard economic methodology prefers non-technical conditions appearing in representation theorems to be formulated directly in terms of the preference relation. This paper provides conditions of this sort: that is, properties of the preference relation which are necessary and sufficient for the existence of a set of acts which can play the role of constant acts and support a state-independent utility representation. Results will be presented both for the Savage framework and the Anscombe and Aumann framework.

To be more precise, state independence comes in two flavours: *ordinal* state independence, according to which the preference order on constant acts is independent of the state, and the stronger notion of *cardinal* state independence, according to which the numerical utilities assigned to the consequences of constant acts are independent of the state. So the initial question can be understood in two ways, depending on which concept of state-dependence is of interest. This paper is mainly concerned with conditions for the existence of a set of acts that can play the role of constant acts and with respect to which there is ordinal state independence; however, the results will yield as a corollary a further condition which, in a large class of cases, guarantees that there is cardinal state independence with respect to this set of acts. Monotonicity – the traditional axiom for ordinal state independence – states that, for any pair of constant acts, the first is preferred to the second if and only if, for any non-null event, the first is preferred to the second given that event (Definition II.3). We say that *essential monotonicity* holds if there is a set of acts which can play the role of constant acts and which satisfies monotonicity (Definition II.4). The main result of this paper will be conditions on the preference relation that are necessary and sufficient for essential monotonicity to hold.

As an illustration, consider the tables in Figure II.1, each representing a simple decision problem where there are three outcomes – c_1, c_2, c_3 – and two states – s_1, s_2; the entries in the respective tables indicate the preference orders on the outcomes conditional on the states.[2] The preference relation displayed in the left-hand table is not state independent. Consider now the set containing the three acts f_1, f_2 and f_3, where $f_1(s_1) = c_1$, $f_1(s_2) = c_3$, $f_2(s_1) = c_2$, $f_2(s_2) = c_2$, $f_3(s_1) = c_3$ and $f_3(s_2) = c_1$. This set of acts does satisfy the monotonicity condition: they lie in the same relation according to the preference order conditional on s_1 as according to the preference relation conditional on s_2. Furthermore, this set of acts can play the role of constant acts, because it has the following property (Section II.2.2): for each state and each outcome, there is a unique act in the set which yields that outcome on that state. Hence, although monotonicity does not hold in this case, essential monotonicity does: there is a set of acts which can play the role of constant acts and which satisfies the monotonicity condition.

By contrast, essential monotonicity does not hold in the example the right-hand table. There is no set of acts that has the property mentioned above and satisfies monotonicity. Such a set of acts must take different values on each state. So, for each such set, there is a

2. We are assuming, for the sake of this example, the basic postulates of Savagean decision theory (weak order and Savage's sure-thing principle, or Anscombe and Aumann's independence), under which there is a well-defined concept of preference order conditional on a state.

Figure II.1 – Essential state independence

	s_1	s_2
c_3	first	third
c_2	second	second
c_1	third	first

<table>
<tr><th></th><th>s_1</th><th>s_3</th></tr>
<tr><td>c_3</td><td>first</td><td>first</td></tr>
<tr><td>c_2</td><td>second</td><td rowspan="2">second equal</td></tr>
<tr><td>c_1</td><td>third</td></tr>
</table>

strict preference between any pair of acts in the set according to the preference order conditional on s_1. However, there exists a pair of acts of the set between which the preference order conditional on s_3 is indifferent (those which take the values c_1 and c_2). Hence, for each set of acts which can play the role of constant acts, there are pairs of acts which lie in different relations according to the preference orders conditional on s_1 and conditional on s_3. Essential monotonicity does not hold.

The problem in this case seems to be that the preference orders conditional on the different states do not have sufficient structure in common for there to exist a set of acts that could play the role of constant acts and satisfy the monotonicity property. The main results of this paper exploit this observation. Monotonicity implies that, to support an ordinally state-independent representation, a set of acts which may count as constant should be such that one of these acts is preferred to another given a particular state if and only if the former is preferred to the latter given any other state. This yields the condition that, for any two states, the order on the outcomes given one state is isomorphic to the order on the outcomes given the other.

This basic idea needs to be qualified. For one, the reasoning only applies to non-null states. This qualification is all that one needs to obtain necessary and sufficient conditions for essential monotonicity when working in the Anscombe and Aumann framework (Theorem II.1, Section II.3). However, in the Savage framework, where all states are null, a further technical condition is required; once again, a set of necessary and sufficient conditions for essential monotonicity can be obtained (Theorem II.3, Section II.4).

A first corollary of the main results are representation theorems, comparable to the one in Karni (1993b). The condition requiring the existence of a set of acts which satisfies monotonicity is replaced by a condition requiring that the conditional preference orders are isomorphic. One thus obtains new representation theorems, in both the Anscombe and

Aumann and Savage frameworks (Theorems II.2 and II.4).

As a second corollary, the main results imply, in a large class of cases, a strong uniqueness property on the sets of acts which satisfies monotonicity and can play the role of constant acts. This yields, in such cases, necessary and sufficient conditions for there to be a set of acts with respect to which there is cardinal state independence.

Section II.2 introduces the decision-theoretic frameworks and the basic notions. Necessary and sufficient conditions for essential monotonicity will be then given in the Anscombe and Aumann framework (Section II.3) and in the Savage framework (Section II.4). The corollaries mentioned above will be drawn. Section II.5 discusses relationship with previous work, notably by Karni Karni (1993b,a), and possible extensions of the results. Proofs of the main results are in the Appendix.

II.2 Preliminaries and Basic Concepts

II.2.1 The two frameworks

The basic notions and results of this paper apply both in the Anscombe and Aumann framework and in the Savage framework. In both cases, S will designate the set of *states of nature* and C the set of *outcomes*.

In the Anscombe and Aumann framework Anscombe and Aumann (1963), S is a finite nonempty set. Subsets of S are called *events*. C is a nonempty finite or countably infinite set; a *consequence* is a probability measure on C with finite support. $\Delta(C)$ is the set of consequences. The set of consequences which assign weight 1 to one outcome is in one-to-one correspondence with the set of outcomes; we denote the former set also by C, with slight abuse of notation. Let H denote the set of mappings $h : S \to \Delta(C)$. In this framework, there is a mixture structure on H, ie. for $h, h' \in H$ and $\alpha \in [0, 1]$, there is $\alpha h + (1 - \alpha)h' \in H$, which is defined by $(\alpha h + (1 - \alpha)h')(s) = \alpha h(s) + (1 - \alpha)h'(s)$.[3] We call the elements of H taking values in C *acts*, to allow simpler comparison with the Savage framework. $\mathcal{A}$ is the set of acts. A preference relation $\preccurlyeq$ is assumed on H; we shall refer to its restriction to $\mathcal{A}$ also by $\preccurlyeq$.

3. Anscombe and Aumann's Reversal of Order axiom (Anscombe and Aumann, 1963, p201) is assumed here.

In the Savage framework Savage (1954), S is an infinite set which is equipped with a σ-algebra $\mathscr{S}$, such that, for any $s \in S$, $\{s\} \in \mathscr{S}$. The elements of $\mathscr{S}$ are called *events*. C is a nonempty finite or countably infinite set. Here, C is the set of *consequences*. An *act* is a function $f : S \to C$ which is measurable with respect to $(S, \mathscr{S})$. $\mathcal{A}$ is the set of acts; a preference relation $\preccurlyeq$ is assumed on it.

For the remainder of Section II.2, everything said will apply to both frameworks.

Throughout this paper, f, g, and so on will designate members of $\mathcal{A}$. As usual, $f \sim g$ expresses indifference according to the preference order, and, for any event A, $\preccurlyeq_A$ denotes the preference order given A.[4] We write $f \preccurlyeq_s g$ instead of $f \preccurlyeq_{\{s\}} g$. For any consequence $c \in C$ (respectively $p \in \Delta(C)$), c will denote the element of $\mathcal{A}$ (resp. H) taking value c (resp. p) for all $s \in S$.

Definition II.1. An event A is *null* if, for any acts f, g such that $f(s) = g(s)$ for any $s \notin A$, $f \sim g$. A state s is said to be null if the event $\{s\}$ is null.

II.2.2 Basis

We wish to consider a generalised monotonicity property involving not constant acts but another set of acts. The set of constant acts has several important properties that any set which can be used instead of it in a generalised monotonicity property should also possess. In particular, for any state and any outcome, there is a unique constant act which, applied to the state, yields the outcome. We call any set of acts with this property a basis.

Definition II.2 (Basis)**.** A *basis* $\mathcal{B}$ is a set of $b^i \in \mathcal{A}$, such that, for each $s \in S$ and for each $c \in C$, there exists $b^i \in \mathcal{B}$ with $b^i(s) = c$ and for all $j \neq i$, $b^j(s) \neq c$.

Let $\mathcal{B}_C$ be the basis of constant acts.

Bases span, in a unique way, the set of all acts: an act, as a measurable function from states to outcomes, can equally be thought of as a measurable function from states to elements of the basis. This is the content of the following proposition.

Proposition II.1. *Consider a basis $\mathcal{B}$. For each $f \in \mathcal{A}$, there is a unique measurable function $f^b : S \to \mathcal{B}$ such that $f(s) = f^b(s)(s)$.*

4. For all $f, g \in \mathcal{A}$, $f \preccurlyeq_A g$ iff, for any $f', g' \in \mathcal{A}$ with $f'(s) = f(s)$ and $g'(s) = g(s)$ for any $s \in A$ and $f'(s) = g'(s)$ for $s \notin A$, $f' \preccurlyeq g'$. Since we assume independence and the sure-thing principle (see Sections II.3 and II.4), this coincides with the ordinary notion of conditional preferences.

Different bases afford different ways of expressing the same acts, just as different bases of a vector space afford different ways of referring to the same vector.

An alternative way of understanding the notion of basis is as follows. For any pair of states s and s', the basis $\mathcal{B}$ associates to any outcome c a unique outcome d such that obtaining c on s "corresponds" to obtaining d on s'. So, for any s and s', $\mathcal{B}$ generates a bijective (ie. one-to-one and onto) mapping $\psi_{ss'} : C \to C$. The family of mappings $\{\psi_{ss'} : s, s' \in S\}$ generated by $\mathcal{B}$ is closed under composition (for any $s, s', s'' \in S$, $\psi_{s's''} \circ \psi_{ss'} = \psi_{ss''}$)[5] and contains the identities (ψ_{ss} is the identity, for each $s \in S$). The converse also holds: for any family Σ of bijective functions from C to itself, indexed by pairs of elements of S, that is closed under composition and contains the identities, there corresponds a unique basis. Formally:

Proposition II.2. *Given a basis $\mathcal{B}$, there is a unique family Σ of bijective functions $\psi_{ss'} : C \to C$ for each $s, s' \in S$ which contains the identities and is closed under composition, such that, for each $s, s' \in S$ and every $b^i \in \mathcal{B}$, $\psi_{ss'}(b^i(s)) = b^i(s')$. Conversely, given such a family Σ of bijective functions containing the identities and closed under composition, there exists a unique basis $\mathcal{B}$ such that, for every pair $s, s' \in S$ and every $b^i \in \mathcal{B}$, $\psi_{ss'}(b^i(s)) = b^i(s')$.*

Thus, instead of thinking in terms of bases, one can equivalently think in terms of families of functions which identify which outcome given one state corresponds to a particular outcome given another state. Whilst the former perspective is adopted in this paper, Karni (1993b,a) use the latter.[6]

II.2.3 Monotonicity

As stated in the Introduction, the axiom with which we are mainly concerned here is monotonicity. Here is a formulation common to the Savage and Anscombe and Aumann frameworks (it is visibly identical to Savage's P3; the relation to Anscombe and Aumann's monotonicity in prizes axiom will be discussed in Section II.3.2).

Definition II.3. *Monotonicity* holds if for every non-null event A, and for all constant acts c and d, $c \preccurlyeq_A d$ iff $c \preccurlyeq d$.

5. In words, this property says that if getting d given s' corresponds to getting c given s, and getting e given s'' corresponds to getting d given s', then getting e given s'' corresponds to getting c given s.

6. Notwithstanding the remark on (Karni, 1993a, p192-193).

Monotonicity refers to the basis of constant acts. The idea that the constant acts may not be "really" constant, but that another set of acts may be instead, leads to the following weakening of the axiom:

Definition II.4. *Essential monotonicity* holds if there exists a basis $\mathcal{B}$ such that, for every non-null event A, and for all $b^i, b^j \in \mathcal{B}$, $b^i \preccurlyeq_A b^j$ iff $b^i \preccurlyeq b^j$. Such a basis is an *essentially monotonic basis*.

Just as monotonicity is an ingredient in traditional representation theorems, essential monotonicity can be used in generalised representation theorems (Sections II.3.2 and II.4.2). However, by contrast with many standard axioms of decision theory, essential monotonicity, as stated, is not an immediate condition on the preference relation. The next section presents a condition formulated directly in terms of the preference relation which is equivalent to essential monotonicity in the Anscombe and Aumann framework and the following section does the same for the Savage framework.

II.3 Characterising essential monotonicity: the Anscombe and Aumann framework

II.3.1 Background assumptions

From the Anscombe and Aumann axioms other than those concerning state-independence (ie. weak order, independence, continuity and reversal of order, but not monotonicity in prizes), it follows that there is a representation of $\preccurlyeq$ by a function $U : S \times C \rightarrow \Re$ such that, for any $h, h' \in H$, $h \preccurlyeq h'$ iff

(II.1) $$\sum_{s \in S} \sum_{c \in supp(h(s))} U(s,c)h(s)(c) \leqslant \sum_{s \in S} \sum_{c \in supp(h(s))} U(s,c)h(s)(c)$$

Further, this function is unique up to similar positive affine transformations (or, to use the terminology in Karni et al. (1983), cardinal unit comparable transformations): U' satisfies (II.1) if and only if there is a positive real number a and real numbers b_s for each $s \in S$ such that $U'(s,c) = aU(s,c) + b_s$ for all $s \in S$, $c \in C$ (Fishburn, 1970, Ch 13).

Throughout this section, it is assumed that $\preccurlyeq$ satisfies conditions sufficient for a representation of the form (II.1).

II.3.2 Ordinal and cardinal state independence

Given the background assumptions, the following is equivalent to Anscombe and Aumann's axiom for state-independence of utility, mononiticity in prizes.

Definition II.5. *Monotonicity in consequences* holds if for every non-null event A, and for all constant-valued $p, q \in H$, $p \preccurlyeq_A q$ iff $p \preccurlyeq q$.

Monotonicity, as defined in Definition II.3, is monotonicity in outcomes (elements of C) rather than monotonicity in consequences (elements of H). Monotonicity in outcomes only ensures *ordinal* state independence, while the stronger monotonicity in consequences is required to ensure *cardinal* state independence.

We thus propose factorising the monotonicity-in-consequences condition into an ordinal state-independence condition – the notion of monotonicity in Definition II.3 – and a condition (to be stated below) which, given ordinal state independence, yields cardinal state independence. Not only does this underline the common ground with the Savage framework, where there are two axioms for state independence (Section II.4.2), but it also allows a more precise delineation of the subject of this paper. We are considering the weakening of ordinal state independence that can be obtained by choosing a new set of constant acts.

The condition for cardinal state independence must presuppose that the generalised condition for ordinal state independence is satisfied; namely, that there is an essentially monotonic basis. In the standard case this is assumed to be the set of constant acts, but in the interests of generality, no assumptions will be made regarding it here, so it has to be explicitly mentioned.

Definition II.6. *Cardinal state independence* holds with respect to an essentially monotonic basis $\mathcal{B}$ if for any $h, h' \in H$ which are mixtures of elements of $\mathcal{B}$, ie. $h = \sum_{i=1}^{n} \alpha_i b^i$ and $h' = \sum_{j=1}^{m} \beta_j b^j$ with $\alpha_i, \beta_j \geqslant 0$, $\sum \alpha_i = \sum \beta_j = 1$, and for any non-null event A, $h \preccurlyeq_A h'$ iff $h \preccurlyeq h'$.

Clearly, the conjunction of essential monotonicity with essentially monotonic basis $\mathcal{B}$ and cardinal state independence with respect to $\mathcal{B}$ is equivalent to a generalised version of monotonicity in consequences, where the implicit use of the basis of constant acts is replaced by the essentially monotonic basis $\mathcal{B}$. Anscombe and Aumann's representation theorem Anscombe and Aumann (1963) can thus be immediately generalised, replacing monotonicity in consequences by the conjunction of essential monotonicity and cardinal state independence with respect to an essentially monotonic basis $\mathcal{B}$, and the representation by a probability and a utility function on the set of outcomes by a representation by a probability and a utility function on the basis $\mathcal{B}$. That is, there is a representation of the following form:

$$\text{(II.2)} \qquad \sum_{s \in S} \sum_{\substack{b^i \in \mathcal{B} \ s.t. \\ b^i(s) \in \\ supp(h(s))}} p(s)u(b^i)h(s)(b^i(s)) \leqslant \sum_{s \in S} \sum_{\substack{b^i \in \mathcal{B} \ s.t. \\ b^i(s) \in \\ supp(h'(s))}} p(s)u(b^i)h'(s)(b^i(s))$$

where p is unique, and u is unique up to positive affine transformations. This generalisation can be thought of as a representation theorem for state-dependent utility, insofar as the equation (II.2) can be reformulated as

$$\sum_{s \in S} \sum_{\substack{c \in C \ s.t. \\ c \in supp(h(s))}} p(s)u'(s,c)h(s)(c) \leqslant \sum_{\substack{c \in C \ s.t. \\ c \in supp(h(s))}} p(s)u'(s,c)h'(s)(c)$$

with $u'(s,c) = u(b^i)$, where $b^i(s) = c$. This was one of the motivations of Karni (1993b,a).

II.3.3 Characterising essential monotonicity

The aim is to find a property of the preference relation which is necessary and sufficient for the existence of an essentially monotonic basis. The major obstacle to the existence of this sort of basis was illustrated in the introduction: if, for non-null states s and t, $\leqslant_s$ and $\leqslant_t$ do not have a sufficient number of properties in common then no essentially monotonic basis exists. Intuitively, it seems evident how much the two orders need to have in common: they should be isomorphic. The main theorem of this section confirms that this is all that is required. First a preliminary definition.

Definition II.7. Two ordered sets (or orders, for short) $(X, \preceq_X)$ and $(Y, \preceq_Y)$ are *isomorphic* if there exists a bijective function $\psi : X \to Y$ such that, for all $x, x' \in X$, $x \preceq_X x'$ iff $\psi(x) \preceq_Y \psi(x')$.

For any state s, let $(C, \leqslant_s)$ be the set C equipped with $\leqslant_s$. The condition that the preference orders conditional on states have the same structure only applies to non-null states.

Definition II.8. *Isomorphic conditional preferences* holds if for any non-null $s_1, s_2 \in S$, $(C, \preccurlyeq_{s_1})$ and $(C, \preccurlyeq_{s_2})$ are isomorphic.

The necessity and sufficiency of this condition is expressed by the following theorem.

Theorem II.1. *Essential monotonicity holds iff isomorphic conditional preferences holds.*

Taken in tandem with the generalisation of Anscombe and Aumann's theorem described above, Theorem II.1 has the following representation theorem as an immediate corollary.

Theorem II.2. *Suppose that there is a representation of $\preccurlyeq$ of sort specified in (II.1) which is unique up to similar positive affine transformations. Then:*

1. *the following are equivalent:*
   - *(i) isomorphic conditional preferences holds and cardinal state independence holds with respect to an essentially monotonic basis $\mathcal{B}$*
   - *(ii) there exists a probability distribution p on S, a basis $\mathcal{B}$ and a real-valued function u on $\mathcal{B}$, such that, for $h, h' \in H$, $h \preccurlyeq h'$ iff*

$$\sum_{s\in S} \sum_{\substack{b^i\in\mathcal{B} \text{ s.t.} \\ b^i(s)\in \\ supp(h(s))}} p(s)u(b^i)h(s)(b^i(s)) \preccurlyeq \sum_{s\in S} \sum_{\substack{b^i\in\mathcal{B} \text{ s.t.} \\ b^i(s)\in \\ supp(h'(s))}} p(s)u(b^i)h'(s)(b^i(s)) \qquad \text{(II.3)}$$

2. *for a given $\mathcal{B}$, p is unique, and u is unique up to positive affine transformations.*

As noted above, this can be thought of as a representation theorem for state-dependent utility.

A corollary of the proof of Theorem II.1 is the following uniqueness property for essentially monotonic bases.

Corollary II.1. *Suppose that essential monotonicity holds, and let $\mathcal{B}_1$ and $\mathcal{B}_2$ be two essentially monotonic bases. Suppose furthermore that, for any non-null s, there is either a minimal or a maximal element (or both) in $(C, \preccurlyeq_s)$. Then, for each $b_1^i \in \mathcal{B}_1$, there is a $b_2^j \in \mathcal{B}_2$ such that, for any non-null event A, $b_1^i \sim_A b_2^j$. Furthermore, for each $b_1^i \in \mathcal{B}_1$, $|\{b_2^j \in \mathcal{B}_2 |$ for any non-null event $A,\ b_1^i \sim_A b_2^j\}| = |\{b_1^k \in \mathcal{B}_1 |\ b_1^i \sim b_1^k\}|$.*

In a word, if $(C, \preccurlyeq_s)$ has a maximal or minimal element, the only differences between essentially monotonic bases occur when there are several outcomes on which $\preccurlyeq_s$ is indifferent. For example, there could be essentially monotonic bases $\mathcal{B}$ and $\widehat{\mathcal{B}}$, with the first containing elements b^i, b^j with $b^i(s_1) = c$, $b^i(s_2) = d$, $b^j(s_1) = c'$, and $b^j(s_2) = d'$ and the second containing elements $\widehat{b^i}$, $\widehat{b^j}$ with $\widehat{b^i}(s_1) = c$, $\widehat{b^i}(s_2) = d'$, $\widehat{b^j}(s_1) = c'$, and $\widehat{b^j}(s_2) = d$, but only if $c \sim_{s_1} c'$ and $d \sim_{s_2} d'$.

The requirement that there is a maximal or minimal element is necessary, as can be seen from the case where the preference orders conditional on states are isomorphic to the order on the integers $((\mathbb{Z}, \preccurlyeq))$. Any basis obtained from a given essentially monotonic basis by "shifting" the values that the elements take on a particular state by a constant number of rungs is also an essentially monotonic basis.

Corollary II.1 says that whenever C is finite or countably infinite with a maximal or minimal element according to $\preccurlyeq_s$, there is a little degree of freedom in the choice of essentially monotonic basis, supposing that essential monotonicity holds. In such cases, there is thus no degree of freedom in whether cardinal state independence holds: cardinal state independence holds with respect to a particular essentially monotonic basis if and only if it holds with respect to *any* essentially monotonic basis. In such cases, Theorem II.2 can be strengthened: "an essentially monotonic basis" in Clause 1. (i) can be replaced with "any essentially monotonic basis". Furthermore, the relativity to the basis in the uniqueness clause (Clause 2.) can be effectively removed, since the same probabilities and utilities (up to positive affine transformations) are obtained for all essentially monotonic bases. This is not so in cases of unbounded $\preccurlyeq_s$. Consider a development of the previous example where the consequences are the set of integers and the utility of a consequence c given any state s is $2.c$ if $c > 0$ and c if not. There is cardinal state independence with respect to the basis of constant acts, $\mathcal{B}_C$; however, for any basis obtained from $\mathcal{B}_C$ by "shifting" in the way described above, although it is essentially monotonic, cardinal state independence does not hold with respect to it.

II.4 Characterising essential monotonicity: the Savage framework

II.4.1 Background assumptions

We want to assume an equivalent of the representation (II.1) for the Savage framework. Since the set of states S is infinite, the double sum must be replaced by an integral over the product space $S \times C$. Accordingly the function U is replaced by a measure on $(S \times C, \mathscr{T})$, where the σ-algebra $\mathscr{T}$ is the the product of $\mathscr{S}$ and the power set of C. Finally, an act $f \in \mathcal{A}$ is replaced by its graph $(\{(s,c)|f(s) = c\})$; since f is a measurable function, its graph is a measurable set in $(S \times C, \mathscr{T})$. (We use the symbol f to refer both to the function and its graph.) We assume here that $\preccurlyeq$ satisfies conditions sufficient for the existence of a measure U on $(S \times C, \mathscr{T})$ such that, for any $f, f' \in \mathcal{A}$, $f \preccurlyeq f'$ iff

$$\int_f dU \leqslant \int_{f'} dU \tag{II.4}$$

and such that this measure is unique up to similar positive affine transformations: U' satisfies (II.4) if and only if there exists $a > 0$ and a measurable function $b : S \to \Re$ such that $U' = aU + b$. Sufficient conditions for this representation have been proposed in Hill (2010): they correspond more or less to the basic Savage axioms except for those concerning state-independence (in particular, weak order, the sure-thing principle and continuity axioms, but not monotonicity and weak comparative probability).

II.4.2 Ordinal and cardinal state independence

As in the Anscombe and Aumann framework, monotonicity (Definition II.3) only ensures *ordinal* state independence, but not *cardinal* state independence (see also Karni (1993b); Karni and Mongin (2000)). The further condition proposed by Savage is his P4 (also known as weak comparative probability), which we shall call *cardinal state independence*, to underline the analogy with the Anscombe and Aumann framework. The traditional formulation assumes that the set of constant acts is an essentially monotonic basis; in the interests of generality, assumptions regarding the appropriate essentially monotonic basis are dropped in the formulation used here.

Definition II.9. *Cardinal state independence* holds with respect to an essentially monotonic basis $\mathcal{B}$ if for every pair of events A and B and every $b^i, b^j, b^k, b^l \in \mathcal{B}$ such that $b^i \prec b^j$ and $b^k \prec b^l$, $b^i_A b^j \preccurlyeq b^i_B b^j$ iff $b^k_A b^l \preccurlyeq b^k_B b^l$, where $b^i_A b^j$ is the act which takes the values of b^i on A and the values of b^j on A^c.

Just as there is an immediate generalisation of Anscombe and Aumann's theorem (Section II.3.2), there is an immediate generalisation of Savage's representation theorem Savage (1954), where his state-independence axioms are replaced by essential monotonicity and cardinal state independence with respect to an essentially monotonic basis $\mathcal{B}$, and these are necessary and sufficient, given the background assumptions, for a representation of the following form:

$$\int_S u(f^b(s))dp \preccurlyeq \int_S u(f'^b(s))dp$$

where f^b and f'^b are as in Proposition II.1 and p and u have the ordinary uniqueness properties. As for the Aumann and Anscombe case, this generalisation of Savage's theorem can be thought of as a representation theorem for state-dependent utility.

II.4.3 Characterising essential monotonicity

The isomorphic conditional preferences condition used in the Anscombe and Aumann framework (Definition II.8) cannot be directly applied in the Savage framework, because in Savage's theory, where the state space is atomless, all states are null, so the condition is trivial. Instead, the existence of essentially monotonic bases will be characterised in terms of preferences conditional on non-null events belonging to a particular set of events. These non-null events play the role that states did in the previous section: the new isomorphic conditional preferences condition states that the preference orders conditional on these non-null events agree. However, for there to be an essentially monotonic basis, the preference orders conditional on *all* non-null events need to agree (Definition II.4). To guarantee this, the non-null events used in the isomorphic conditional preferences condition need to have the appropriate stability property: the preference order conditional on any subevent must coincide with the preference order conditional on the event itself.

Definition II.10. A non-null event E is *stable for* $c, d \in C$ if the following condition holds: $c \preccurlyeq_E d$ iff for every non-null event A such that $A \subseteq E$, $c \preccurlyeq_A d$. E is said to be *stable* if it is stable for all $c, d \in C$.

The following technical condition guarantees that an appropriate set of stable events exists.

Definition II.11. *Local stability* holds if, for any $f, g \in \mathcal{A}$, there exists events E_1, E_2 and E_3, such that the non-empty E_i form a partition of S and such that: for any non-null event $A \subseteq E_1$, $f \prec_A g$; for any non-null event $A \subseteq E_2$, $f \sim_A g$; and for any non-null event $A \subseteq E_3$, $f \succ_A g$.

Proposition II.3. *Suppose that local stability holds. Then there exists a partition $\mathcal{P}$ of S such that each non-null $E \in \mathcal{P}$ is stable.*

The partition in this proposition is exactly of the sort needed to characterise essential monotonicity. Where the condition proposed in the previous section concerns preference orders conditional on non-null states, the condition relevant here involves preference orders conditional on the non-null elements of the partition of stable events.

Definition II.12. Assume that local stability holds and let $\mathcal{P}$ be a partition of S such that each $E \in \mathcal{P}$ is stable. Then *isomorphic conditional preferences* holds if for any non-null $E, E' \in \mathcal{P}$, $(C, \preccurlyeq_E)$ and $(C, \preccurlyeq_{E'})$ are isomorphic.[7]

It turns out that local stability is implied by essential monotonicity, so we have the following necessary and sufficient conditions for essential monotonicity in the Savage framework.

Theorem II.3. *Essential monotonicity holds iff local stability and isomorphic conditional preferences hold.*

As in the Anscombe and Aumann case, this has a representation theorem as an immediate corollary.

Theorem II.4. *Suppose that there is a representation of $\preccurlyeq$ of the sort specified in (II.4) which is unique up to similar positive affine transformations. Then:*

7. Given that, for any pair of partitions with stable elements, each element of one will have a non-null overlap with at least one element of the other, isomorphic conditional preferences holds independently of the partition chosen, and so is well-defined.

1. *the following are equivalent:*

   (i) local stability and isomorphic conditional preferences hold and cardinal state independence holds with respect to an essentially monotonic basis $\mathcal{B}$

   (ii) there exists a probability distribution p on S, a basis $\mathcal{B}$ and a real-valued function u on $\mathcal{B}$, such that, for $f, f' \in \mathcal{A}$, $f \preccurlyeq f'$ iff

(II.5)
$$\int_S u(f^b(s))dp \leqslant \int_S u(f'^b(s))dp$$

   where f^b and f'^b are as in Proposition II.1.

2. *for a given $\mathcal{B}$, p is unique, and u is unique up to positive affine transformations.*

A version of Corollary II.1 continues to hold in the Savage framework, with mention of states s replaced by elements E of a partition of S consisting entirely of stable events. The remarks concerning Theorem II.2 and Corollary II.1 also apply here.

II.5 Discussion

As mentioned above, the idea of replacing the set of constant acts by a set of "really" constant acts has already been employed by Edi Karni, both in the Anscombe and Aumann framework Karni (1993a) and in the Savage framework Karni (1993b). Leaving aside the technical differences between this paper and Karni's – Karni employs state-specific transformations of outcomes, whereas we use bases (Proposition II.2 states that they are equivalent), Karni considers equivalence classes of outcomes whereas until now we have considered the outcomes themselves (but see below) [8] – there are fundamental differences between the goals of Karni's papers and this one.

In his work on the Anscombe and Aumann framework, Karni makes assumptions which are essential to his theorem and which imply that the set of outcomes is uncountably infinite; [9] by contrast, this paper only considers the finite and countably infinite cases. Moreover, his focus seems to be cardinal state independence, whereas the main concern here is

8. (Karni, 1993a, p191) explicitly states this. Karni (1993b) uses transformation functions that are onto though not necessarily one-to-one; these can be seen as bijective functions on equivalence classes.

9. He assumes the existence of a "conditional certainty equivalent" outcome for every consequence (p191).

ordinal state independence. Indeed, in the main result of his paper, he assumes monotonicity (Definition II.3), and then constructs a basis which would be capable of supporting a state-independent representation; in the presence of monotonicity, the condition he places on this basis – state-invariance – corresponds to what we have called cardinal state independence (Definition II.6).[10] These differences are perhaps not unrelated. When the set of outcomes is finite or countably infinite, the problem of finding a cardinally state independent basis is trivial in many cases, once the problem of finding an ordinally state independent basis has been solved (Corollary II.1); ordinal state dependence thus seems a more natural subject of concern in the finite and countably infinite cases.

In the Savage case Karni (1993b), Karni lays down some conditions on a basis; his main theorem states that, if there exists a basis satisfying these conditions, then there is a representation similar to (II.5).The main contrast between Karni's result and the representation theorem in Section II.4 (Theorem II.4) is the use of a condition requiring the existence of a particular basis in the former and the use of a condition requiring isomorphism of conditional preferences in the latter. From the viewpoint of standard economic methodology, there are two significant differences between these conditions.

Firstly, it is not immediately clear what property of the preference relation is demanded by the condition requiring the existence of a particular basis: it says that there is a set of acts such that the preferences over these acts have certain properties, but it does not specify which preference relations admit the existence of such sets of acts. By contrast, the condition on conditional preferences is formulated directly as a property of the preference relation: it says that the preference order given one stable event is the same as the preference order given any other stable event (in the sense of "the same" which is relevant for orders). This difference is significant if one supposes that the preference comparisons are observable from choices: it is not obvious what constraints Karni's condition implies on the agent's behaviour, whereas it can be straightforwardly checked, by comparing the agent's preferences conditional on different events, whether the condition proposed here is satisfied or not.

Secondly, the condition requiring the existence of a particular basis involves existential quantification over the objects of the theory: it quantifies across sets of acts, and the theory

10. He notes in the final section of his paper that the monotonicity assumption is not required for his theorem.

is dealing with preferences over a given set of acts. By contrast, the condition requiring isomorphic conditional preferences does not involve existential quantification over objects of the theory. At most, it involves existential quantification over technical constructions, insofar as the definition of isomorphism given above (Definition II.7) involves quantification over mappings between orders. However, the notion of isomorphism can in fact be defined without using existential quantification at all, if one employs the machinery of modern logic: two orders are isomorphic if everything you can say about the first (in an appropriately formalised language) holds for the second. [11] This difference has consequences for how the conditions are to be understood. Conditions which involve existential quantification over objects of the theory are usually thought of as "technical" or "structural": continuity axioms, which normally involve existential quantification over acts or events, are standard examples. Therefore, the condition requiring the existence of an appropriate basis is most naturally thought of as a technical condition, whereas the condition requiring isomorphic conditional preferences is not. As far as standard economic methodology is concerned, the result in Section II.4 differs from that in Karni (1993b) with respect to the possibility of applying the central conditions to agents' behaviour, and the interpretation of these conditions as technical or not.

Let us conclude by noting two extensions of the proposal made here. Firstly, since nothing in the isomorphic conditional preferences condition demands that the orders which are isomorphic are orders on the same set C, the results apply if it is assumed that there is a set C_s of outcomes for each state s. [12] If the orders $(C_s, \preccurlyeq_s)$ are isomorphic (and thus have the same cardinality), one can define a basis which "ties together" the different sets of outcomes and is essentially monotonic, so that, if cardinal state-independence holds as well, there is a representation of the form (II.3). A special case is when C_s is taken to be the set of equivalence classes of C under $\sim_s$. There may be cases where isomorphic conditional preferences does not hold for the initial set of outcomes, but does hold on the equivalence classes. The results above can be applied, yielding a representation of preferences over functions from states to equivalence classes of outcomes (or lotteries over equivalence classes of outcomes), which can be translated back into a representation of

11. An earlier version of this paper presented the result in this form. See any logic textbook, such as Chang and Keisler (1990), for details.

12. This is the sort of setup proposed in (Hammond, 1998, §6), for example. In the text, the Anscombe and Aumann framework is discussed, but similar points hold for the Savage framework.

preferences over acts. This is the technique used in Karni Karni (1993b,a).

Secondly, other results may be extended by weakening the monotonicity axiom in much the same way that Savage's and Anscombe and Aumann's theorems have been extended here. For example, several state-dependent utility theories have been proposed which assume monotonicity but drop the cardinal state independence axiom Karni and Schmeidler (1993); Wakker and Zank (1999). The characterisation theorems proved here provide an immediate generalisation of these results: monotonicity is replaced by isomorphic conditional preferences (and, in the Savage framework, local stability), and any weakened version of the cardinal state independence axiom is reformulated in terms of essentially monotonic bases.

II.A Appendix

Proof of Proposition II.1. Set $f^b(s) = b^i$ where b^i is the unique element of the $\mathcal{B}$ such that $b^i(s) = f(s)$. It remains to show that f^b is measurable. For each i and for each $c \in C$, $\{s \mid f(s) = c = b^i(s)\} = \{s \mid f(s) = c\} \cap \{s \mid b^i(s) = c\}$ and hence is measurable. For each i, $\{s \mid f(s) = b^i(s)\} = \bigcup_{c \in C} \{s \mid f(s) = c = b^i(s)\}$ is a countable union of measurable sets and thus measurable. Hence f^b is measurable.

□

Proof of Proposition II.2. Let $\mathcal{B}$ be a basis. For $s, s' \in S$, define $\psi_{ss'} : C \to C$ as follows: for each $c \in C$, take the unique $b^i \in \mathcal{B}$ such that $b^i(s) = c$ and set $\psi_{ss'}(c) = b^i(s')$. It is readily checked that this is a well-defined bijective function and that the set $\Sigma = \{\psi_{ss'} \mid s, s' \in S\}$ has the required properties.

Let Σ be a family of the sort required. Pick any $t \in S$. Define the basis $\mathcal{B}$ as follows: the elements b^c are indexed by C, and $b^c(s) = \psi_{ts}(c)$. Since $\psi_{ss'}$ are bijections, $\mathcal{B}$ is a basis; moreover, it is readily checked that it has the required properties.

□

Proof of Proposition II.3. For each $c_j, c_k \in C$, there exists by local stability a partition into at most three elements which are stable for c_j, c_k. Call the partition $\mathcal{P}^{c_j c_k}$.

Let $\mathcal{Q}$ be the coarsest common refinement of $\mathcal{P}^{c_j c_k}$ for all $c_j, c_k \in C$.[13] Since C is at most countable, each element of $\mathcal{Q}$ is the intersection of at most countably many events; so it is an event. Furthermore, by construction, each non-null element of $\mathcal{Q}$ is a subevent of an event of $\mathcal{P}^{c_j c_k}$, for each $c_j, c_k \in C$. Since each element of $\mathcal{P}^{c_j c_k}$ is stable for c_j and c_k, each element of $\mathcal{Q}$ is also stable for c_j and c_k. This holds for all $c_j, c_k \in C$, hence every non-null element of $\mathcal{Q}$ is stable for all $c_j, c_k \in C$.

□

The following lemma shall be useful in the proof of Theorem II.3.

Lemma II.A.1. *Suppose the Savage framework, so that S is infinite. If essential monotonicity holds, then local stability holds.*

13. That is, the coarsest partition each of whose elements are contained in a single element of $\mathcal{P}^{c_j c_k}$ for each c_j, c_k.

Proof. Let $f, g \in \mathcal{A}$. By Proposition II.1, there are corresponding measurable $f^b, g^b : S \to \mathcal{B}$; hence there are partitions $\mathcal{P}^f$, $\mathcal{P}^g$ of S such that f^b (respectively, g^b) takes constant values on each of the elements of $\mathcal{P}^f$ (resp. $\mathcal{P}^g$). Let $\mathcal{Q}$ be the coarsest common refinement of $\mathcal{P}^f$ and $\mathcal{P}^g$. On each element E of $\mathcal{Q}$, both f^b and g^b take constant values in $\mathcal{B}$; hence, since $\mathcal{B}$ is an essentially monotonic basis, $f \preccurlyeq_E g$ iff, for any non-null event $A \subseteq E$, $f \preccurlyeq_A g$. The partition required for local stability is readily constructed from $\mathcal{Q}$.

□

The following lemma is at the heart of Theorems II.1 and II.3.

Lemma II.A.2. *Suppose either the Ancombe and Aumann or the Savage framework, and suppose that there exists a partition $\mathcal{P}$ of S such that each non-null element of $\mathcal{P}$ is stable. Then essential monotonicity holds iff $(C, \preccurlyeq_E)$ and $(C, \preccurlyeq_{E'})$ are isomorphic for all non-null $E, E' \in \mathcal{P}$.*

Proof. Left to right. Suppose that essential monotonicity holds. Then there is a basis $\mathcal{B}$ such that, for any non-null event A and all $b^i, b^j \in \mathcal{B}$, $b^i \preccurlyeq b^j$ iff $b^i \preccurlyeq_A b^j$. So, for each non-null $E, E' \in \mathcal{P}$ and all $b^i, b^j \in \mathcal{B}$, $b^i \preccurlyeq_E b^j$ iff $b^i \preccurlyeq_{E'} b^j$. The identity map on $\mathcal{B}$ is an isomorphism between $(\mathcal{B}, \preccurlyeq_E)$ and $(\mathcal{B}, \preccurlyeq_{E'})$.[14] We show that, for any $E \in \mathcal{P}$, $(\mathcal{B}, \preccurlyeq_E)$ and $(C, \preccurlyeq_E)$ are isomorphic. This will be sufficient to prove the result, because it follows that, for any $E, E' \in \mathcal{P}$, $(C, \preccurlyeq_E)$ is isomorphic to $(\mathcal{B}, \preccurlyeq_E)$, which is isomorphic to $(\mathcal{B}, \preccurlyeq_{E'})$, which is isomorphic to $(C, \preccurlyeq_{E'})$.

Fix $E \in \mathcal{P}$. Since each $b^i \in \mathcal{B}$ is measurable, the set $\{s \in E | \, b^i(s) = c\}$ is measurable, for all $c \in C$. So, to each b^i, there is an associated partition of E, $\mathcal{P}^{b^i} = \{\{s \in E | \, b^i(s) = c\} | \, c \in C\}$. Since C is at most countable, $\mathcal{B}$ is, and thus the coarsest common refinement of the $\mathcal{P}^{b^i}$ exists; call it $\mathcal{Q}$. By construction, for each non-null $A \in \mathcal{Q}$ and each $b^i \in \mathcal{B}$, $b^i(s) = b^i(s')$ for all $s, s' \in A$. So $\psi_A : \mathcal{B} \to C$, where $\psi_A(b^i) = b^i(s)$ for any $s \in A$, is a well-defined function. Since $\mathcal{B}$ is a basis, ψ_A is bijective; moreover, it trivially preserves the order $\preccurlyeq_A$. So, for each non-null $A \in \mathcal{Q}$, $(\mathcal{B}, \preccurlyeq_A)$ and $(C, \preccurlyeq_A)$ are isomorphic. Since $\mathcal{B}$ is a basis, for each non-null $A \in \mathcal{Q}$, $(\mathcal{B}, \preccurlyeq_A)$ is isomorphic to $(\mathcal{B}, \preccurlyeq_E)$; moreover, since E is stable, $(C, \preccurlyeq_A)$ is isomorphic to $(C, \preccurlyeq_E)$. Hence, $(\mathcal{B}, \preccurlyeq_E)$ is isomorphic to $(C, \preccurlyeq_E)$.

Right to left. Suppose that $(C, \preccurlyeq_E)$ to $(C, \preccurlyeq_{E'})$ are isomorphic for any $E, E' \in \mathcal{P}$. There thus exists a family Σ' of bijective order-preserving functions $\psi_{EE'} : C \to C$ for

14. Just as for the case of $(C, \preccurlyeq)$, $(\mathcal{B}, \preccurlyeq')$ is the set $\mathcal{B}$ equipped with the order $\preccurlyeq'$.

each $E, E' \in \mathcal{P}$ which is closed under composition and contains the identities. Construct a family Σ of bijective functions $\psi_{ss'} : C \to C$, for all $s, s' \in S$, as follows. If $s, s' \in E$ for some $E \in \mathcal{P}$, $\psi_{ss'}$ is the identity. If $s \in E$, $s' \in E'$ for some $E, E' \in \mathcal{P}$, $\psi_{ss'} = \psi_{EE'}$. It follows immediately that Σ is closed under composition and contains the identities. By Proposition II.2, there is a unique basis $\mathcal{B}$ corresponding to Σ. It remains to show that $\mathcal{B}$ is an essentially monotonic basis.

For any non-null $E, E' \in \mathcal{P}$ and every $b^i, b^j \in \mathcal{B}$, $b^i \preccurlyeq_E b^j$ iff $b^i(s) \preccurlyeq_E b^j(s)$ for any $s \in E$ (since b^i, b^j are constant on E) iff $\psi_{EE'}(b^i(s)) \preccurlyeq_{E'} \psi_{EE'}(b^j(s))$ (since $\psi_{EE'}$ preserves order) iff $b^i(s') \preccurlyeq_{E'} b^j(s')$ for any $s' \in E'$ (since $\psi_{ss'} = \psi_{EE'}$ and $\psi_{ss'}(b^i(s)) = b^i(s')$) iff $b^i \preccurlyeq_{E'} b^j$. So, for each $b^i, b^j \in \mathcal{B}$, either $b^i \preccurlyeq_E b^j$ for all $E \in \mathcal{P}$ or $b^i \succ_E b^j$ for all $E \in \mathcal{P}$. By (Savage, 1954, Ch 2, Theorem 2), it follows, in the former (respectively, latter) case, that $b^i \preccurlyeq b^j$ (resp. $b^i \succ b^j$);[15] in other words, $b^i \preccurlyeq_E b^j$ for all $E \in \mathcal{P}$ iff $b^i \preccurlyeq b^j$. It remains to show that this is the case for any non-null event A. Since the events of $\mathcal{P}$ are stable, for every $E \in \mathcal{P}$, $b^i \preccurlyeq_E b^j$ iff $b^i \preccurlyeq_{E \cap A} b^j$ when $E \cap A$ is non-null. So, $b^i \preccurlyeq_E b^j$ for all $E \in \mathcal{P}$ iff $b^i \preccurlyeq_{E \cap A} b^j$ for all non-null $E \cap A$ with $E \in \mathcal{P}$. Applying (Savage, 1954, Ch 2, Theorem 2) again, we have $b^i \preccurlyeq b^j$ iff $b^i \preccurlyeq_A b^j$; so $\mathcal{B}$ is an essentially monotonic basis.

□

Remark II.1 (Corollary II.1). The proof of the left-to-right direction has an immediate corollary that, for any two essentially monotonic bases $\mathcal{B}_1$ and $\mathcal{B}_2$ and for any non-null event E, $(\mathcal{B}_1, \preccurlyeq_E)$ and $(\mathcal{B}_2, \preccurlyeq_E)$ are isomorphic. If the order has a maximal or minimal element, it follows that, for each $b_1^i \in \mathcal{B}_1$, there exists $b_2^j \in \mathcal{B}_2$, such that, for any non-null event E, $b_1^i \sim_E b_2^j$. This implies Corollary II.1.

Proof of Theorem II.1. Immediate from Lemma II.A.2, noting that, in the Anscombe and Aumann framework, the partition into singleton events has the required properties. □

Proof of Theorem II.3. Suppose that essential monotonicity holds. By Lemma II.A.1, local stability holds. Hence, by Proposition II.3, there exists a partition of the sort required for Lemma II.A.2 to hold; by the lemma, isomorphic conditional preferences holds. Suppose now that local stability and isomorphic conditional preferences hold. By Proposition II.3, a

15. Although the theorem is stated by Savage, it is readily seen to apply in the Anscombe and Aumann framework.

partition of the sort required for Lemma II.A.2 exists; by the lemma, essential monotonicity holds.

□

Bibliography

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Chang, C. C. and Keisler, H. J. (1990). *Model Theory*. Elsevier Science Publishing Company, Amsterdam.

Drèze, J. H. (1987). *Essays on Economic Decisions under Uncertainty*. Cambridge University Press, Cambridge.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Hammond, P. J. (1998). Subjective Expected Utility. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, volume 1. Kluwer, Dordrecht.

Hill, B. (2010). An additively separable representation in the Savage framework. *J. Econ. Theory*, 145(5):2044–2054.

Karni, E. (1993a). A Definition of Subjective Probabilities with State-Dependent Preferences. *Econometrica*, 61:187–198.

Karni, E. (1993b). Subjective expected utility theory with state dependent preferences. *J. Econ. Theory*, 60:428–438.

Karni, E. and Mongin, P. (2000). On the Determination of Subjective Probability by Choices. *Management Science*, 46:233–248.

Karni, E. and Schmeidler, D. (1993). On the uniqueness of subjective probabilities. *Economic Theory*, 3:267–277.

Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51:1021–1032.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Schervish, M. J., Seidenfeld, T., Kadane, J. B., Association, S., and Sep, N. (1990). State-Dependent Utilities. *Journal of the American Statistical Association*, 85(411):840–847.

Wakker, P. and Zank, H. (1999). State Dependent Expected Utility for Savage's State Space. *Mathematics of Operations Research*, 24:8–34.

III Living without state-independence of utilities

Abstract

This paper is concerned with the representation of preferences which do not satisfy the ordinary axioms for state-independent utilities. After suggesting reasons for not being satisfied with solutions involving state-dependent utilities, an alternative representation shall be proposed involving state-independent utilities and a *situation-dependent factor*. The latter captures the interdependencies between states and consequences. Two sets of axioms are proposed, each permitting the derivation of subjective probabilities, state-independent utilities, and a situation-dependent factor, and each operating in a different framework. The first framework involves the concept of a *decision situation* – consisting of a set of states, a set of consequences and a preference relation on acts; the probabilities, utilities and situation-dependent factor are elicited by referring to other, appropriate decision situations. The second framework, which is technically related, operates in a fixed decision situation; particular "subsituations" are employed in the derivation of the representation. Possible interpretations of the situation-dependent factor and the notion of situation are discussed. [1]

Keywords: Elicitation; Subjective Probability; Subjective Expected Utility; State-dependent utility; Small worlds.

JEL classification: D81, C60.

1. This is slightly modified version of a paper of the same name that appeared in Theory and Decision, 67(4):405-432, October 2009.

III.1 Introduction

It has been known for many years that the axioms implying the state-independence of utilities which appear in many classical theories of decision under uncertainty, such as those of Savage (1954) and Anscombe and Aumann (1963), do not hold in many important situations, such as cases of life insurance or health insurance (Arrow, 1974; Cook and Graham, 1977). The conclusion normally drawn is that a representation involving *state-dependent* utilities is required; thus the work on this subject (for example Karni et al. (1983); Karni and Mongin (2000); Karni and Schmeidler (1993); Karni (2011); Drèze (1987); Drèze and Rustichini (2004)). However, this is not the only option; in this article, an alternative representation shall be proposed, and a representation theorem for it proved.

To introduce the problems posed by violations of the traditional state-independence axioms, and the motivation for representation theorems for state-dependent utility, consider an example closely related to that proposed by Aumann in his correspondence with Savage (reproduced in Drèze (1987, pp76-81); this version is taken from Karni (1996)). A woman is to undergo a potentially fatal operation, with 50-50 chances of survival. Her husband, who knows the odds, is offered two bets: one pays $100 if the operation is successful and nothing if not; the other pays $100 if the operation is not successful and nothing if it is. A common representation of the bets is given in Figure III.1. They are considered to be acts: functions from a set of states – taken to be success and failure of the operation – to a set of consequences (elements to which the husband allocates utility values) – taken to be winning $100 and winning $0. Although the chances of winning the money are equal in both cases, the husband would apparently choose the former bet, because the $100 would be more precious to him if his wife were with him to enjoy it.

Such behaviour may pose two distinct challenges for state-independent utility theories. First of all, it may cast doubt on the accuracy of such theories as representations of *behaviour*. Secondly, it may cast doubt on the claim that such theories *elicit* the agent's beliefs and desires, in the form of probabilities and utilities. The second challenge is more widespread and more serious than the first.

Consider, firstly, the extent of the challenges. There are many cases where agents who make the choices described in the example have preferences which do not satisfy some axioms of the standard theory – most notably, cases where the axioms ensuring state-

Figure III.1 – 2-consequence formulation of decision problem

		Success	Failure
A	Bet on success	$100	$0
B	Bet on failure	$0	$100

independence of utility (P3 and P4 in Savage (1954), Monotonicity in Anscombe and Aumann (1963); see Section III.2.1) are violated. Such cases pose problems for the behavioural and the elicitation claims of the theory: on the one hand, these are cases where the theory is shown to be behaviourally inaccurate; on the other hand, since the axioms need to be satisfied for the theorems to apply, and for the probability and utility to be elicited using them, these are cases where the theory cannot be used to elicit the agent's attitudes.

However, although Aumann and Savage assumed in their discussion that the state-independence axioms are violated, the pattern of choices is not in fact incompatible with satisfaction of these axioms. Karni (1996) gives a development of Aumann's example where all the standard Savage axioms are satisfied, so that probabilities and utilities can be elicited using Savage's theorem. The probability function elicited assigns more weight to the success of the operation than to its failure; given that the utilities are state-independent, this is the only way to account for the husband's strict preference over the bets. Such examples do not pose a particular challenge for Savage's theory on the behavioural front: to the extent that all the axioms are satisfied, the theory adequately describes the agent's behaviour. However, as Karni points out, they seem to pose a challenge to the theory's pretension to elicit the agent's beliefs (and desires). The probability function elicited allocates different probabilities to the success and failure of the operation, and this contradicts the intuition that the husband believes success to be as likely as failure – an intuition which is sustained not only by the assumption that he is fully cognizant of the odds of survival, but also by some of his other actions (for example, he advises a friend, who is offered the same pair of bets, that the odds of winning are the same). Karni concludes that the probabilities furnished by Savage's theorem do not properly represent the husband's beliefs. This is a case where Savage's theory applies and provides a probability and utility function, but it is doubtful whether these accurately represent the agent's attitudes.

The behavioural problem is also less serious than the elicitation problem. Strictly speaking, one could reply to the former problem by resorting to a representation which does not rely on the state-independence axioms. In the case of expected utility, one obtains an additively separable representation, that is, a representation in terms of a real-valued function on pairs of states and consequences – which we shall call the *evaluation function* [2] – that is of the following form:

(III.1) $$f \preceq g \text{ iff } \sum_{s \in S} U(s, f(s)) \leqslant \sum_{s \in S} U(s, g(s))$$

At a push, one could decompose the evaluation function U into a probability and state-dependent utility ($U(s, x) = p(s).u(s, x)$). However, such a decomposition is entirely arbitrary: any (reasonable) probability function could be used, for example. Representations such as this, perhaps with arbitrary decompositions, are sufficient for many economic applications (though not all: see Karni and Schmeidler (1993, §4), for example). However, they do not provide a viable solution to the elicitation problem: here, one needs a (preferably principled) way to decompose the evaluation function U into a *unique* probability function and a (suitably) *unique* utility function.

As regards this problem, which is doubtless the main challenge posed by the failure of state-independence axioms, two strategies have been proposed. The first is promoted by theorists working on state-dependent utilities. The intuition is that, in the example above, the utility of the consequences *depends* on the states: $100 is more desirable if the wife survives than if she does not. Thus the representation should feature state-dependent utilities. Many state-dependent utility theorems are supposed to provide an answer to the elicitation question: they yield a *unique* probability and a (suitably) *unique* state-dependent utility function, the uniqueness being necessary for the claim that these functions represent the agent's beliefs and utilities (Karni et al., 1983; Karni and Schmeidler, 1993; Karni, 1993a,b, 2007, 2011).

The second strategy is that proposed by Savage in his reply to Aumann. He suggests representing the decision by the same set of states, but using as consequences the following four elements: $100 with survival, $100 with demise, $0 with survival, $0 with demise (Figure III.2). He then evokes the "make-believe" situation where the husband considers that any act (function from states to consequences) taking values in this set of consequences

2. Thanks to Mark Machina and Edi Karni for suggesting this term.

Figure III.2 – 4-consequence formulation of decision problem

		Success	Failure
A	Bet on success	\$100 with wife alive	\$0 without wife
B	Bet on failure	\$0 with wife alive	\$100 without wife
C	Bet on success (but wife murdered)	\$100 without wife	\$0 without wife
D	Bet on failure (but wife resuscitated)	\$0 with wife alive	\$100 with wife alive
E	Bet on failure (but wife murdered or resuscitated)	\$0 without wife	\$100 with wife alive

is an available option – and in particular, the acts which, in the case where the operations fails, would yield \$100 and the wife returned to him in good health (acts D and E). In this situation, the axioms of state-independent utility apply and the classical representation theorems can be employed. These yield, according to a proponent of this strategy, the agent's beliefs and utilities (it is easy to show, in particular, that the problem raised by Karni's aforementioned development of the example no longer applies).

Each of these strategies has its price. Under the state-dependent utility approach, the agent's utilities for the consequences depend on the states involved: it thus becomes impossible to use these utilities in decision situations where the same set of consequences are on offer, but the states of the world are different. Consider the case where the agent is offered a bet on a horse race, with the consequences being appropriate combinations of \$100 and \$0: wouldn't one expect him to have the same utilities in this situation as in the one described in the story above? Under the state-dependent analysis, this cannot be the case, since the states on which the utilities are dependent (success and failure of the operation) are not involved. This naturally casts doubt on whether these theorems are yielding his "real" utilities, for one would expect him to have the same utilities for the same consequences considered in different situations (when betting on his wife's operation, when betting on a horse race, when making investment decisions, when taking out life insurance and so on), and this is not a property of the utilities that the state-dependent utility theorems provide.[3]

To understand the weakness of Savage's strategy, note that it operates in two stages. In the first stage, one rewrites the decision problem with an enlarged set of consequences; this corresponds to moving from Figure III.1 to Figure III.2. In the second stage, one asks the decision maker to make his choice *as if* a certain number of the acts – such as acts involving murder or resuscitation (acts C, D and E in Figure III.2) – were not fantastic or ridiculous. The first stage of the strategy involves a simple change in the formulation of the decision problem and no behavioural change on the part of the agent. By contrast, the second stage involves an explicit request that the agent alter his behaviour. This can be seen by the fact that the agent, answering in the "as if" mode, is indifferent between the the acts A and E in Figure III.2, whereas in reality he would never consider picking act E.

In the terminology which will be introduced in the next section, we will say that the second stage involves two different decision situations: both the decision situations have the same states and consequences (and thus acts), but the decision maker's preferences differ between the two. One of the decision situations is the one the agent is actually in (those are the preferences he actually has); the challenge is to understand his behaviour and attitudes in this situation. The other decision situation is "make-believe" – the preferences he has in that situation are hypothetical, and conflict with his real preferences. This is the situation where one can elicit his beliefs and utilities using standard state-independent utility results.

Although these are behaviourally distinct decision situations, for the beliefs and utilities in the former situation to be elicited in the way that Savage suggests, it needs to be assumed that the agent has the same attitudes in the two situations. What permits this assumption? An adequate answer to this question should make clear the relationship between the two situations. The simplest and most intuitive response is that the former situation is as the latter, with an added restriction on the set of acts available (namely, that the fantastic or ridiculous acts are not available). However, the status of this restriction is problematic. First of all, it is posed exogenously to the decision theory used. Furthermore, it is known that one cannot incorporate such restrictions endogenously into representation theorems such as Savage's, since these depend on the availability of a full set of acts (that is, all functions from states to consequences). Indeed, such restrictions generally lead to representations

3. See Section III.3 for further discussion of the relationship between attitudes in different situations.

in terms of state-dependent utilities, which is exactly what Savage's strategy was meant to avoid (Hammond (1998, §6) considers the case, pertinent for this example, where the consequences attainable from different states are distinct). To sum up, the strategy proposed by Savage relies on a decision situation where the agent shows preferences which differ from those he actually has; the weakness in the strategy is that, although the strategy relies on assumptions about the relationship between this situation and the situation the decision maker is actually in (his real preferences), this relationship is poorly understood. Indeed, any attempt to account for the relationship in the theory leads us back to state-dependent utilities.

Two morals can be drawn from these considerations. The discussion of Savage's strategy makes it clear that there is an (inevitable) interdependence between states and consequences in situations such as the one in the example. However, given what has been said about the solution offered by state-dependent utility theorems, it is not clear that this interdependence should be built into the utilities. This seems to suggest that, if there is interdependence of states and consequences, then it must be specific to situations where these states and consequences are involved. A natural suggestion is to represent the interdependence by a *situation-dependent factor* expressing the relationship between states and consequences. The representation would then be as follows. For S the set of states involved in the situation, C the set of consequences and $\preceq$ a preference relation on acts (functions from S to C) there is a probability function p on S, a (state-independent) utility function u on C and a function γ on $S \times C$ – the situation-dependent factor – such that, for any acts f, g,

(III.2) $$f \preceq g \text{ iff } \sum_S p(s).\gamma(s, f(s)).u(f(s)) \leqslant \sum_S p(s).\gamma(s, g(s)).u(g(s))$$

The main aim of this paper is to propose a representation theorem for this sort of representation: that is, to propose a set of conditions under which there is a *unique* probability and a (suitably) *unique* utility and situation-dependent factor such that (III.2) holds. This can be considered as another reply to the behavioural problem posed by examples such as Aumann's, insofar as it represents correctly the preferences of the agent, using not only the probabilities and utilities, but also the situation-dependent factor. It constitutes a reply to the elicitation problem, insofar as it yields (suitably) unique probabilities and utilities, which feature in the representation of preferences and can be thought of as capturing the agent's beliefs and desires. Indeed, the representation theorem relies on an elementary but attractive intuition on how to elicit probabilities and utilities: for elicitation, use situations which are as simple as possible. To elicit the probabilities and utilities in a situation where there is interdependence between states and consequences, the technique will be to look in situations where either the sets of states or sets of consequences are different, but where there is no interdependence between the states and the consequences, so that the traditional (state-independence) results can be applied.

In Section III.2, a formal notion of situation shall be defined, postulates and axioms shall be proposed and a representation theorem shall be stated. In Section III.3, the concepts and techniques introduced, as well as possible interpretations, applications and relations to existing approaches, shall be considered at length. It turns out that the use of other decisions situations is, to a large extent, an artifice of the way the decision problem is modelled; in Section III.4, a representation theorem for (III.2) which operates in a single situation shall be stated. Proofs are to be found in the Appendix.

III.2 Situational version

III.2.1 Preliminaries and axioms

A decision situation is characterised by a set of states and a set of consequences, with a preference on the acts.

Definition III.1. A *decision situation* σ consists of a set S_σ (of states in the situation), a set C_σ (of consequences in the situation) and a binary relation $\preceq_\sigma$ on the set of functions $\mathcal{A}_\sigma$ from S_σ into C_σ (preferences in the situation over the acts in the situation).

A set of situations $\mathcal{S}$ is assumed to be given; this is the set of decision problems which the agent might have faced at the current moment, with his preferences over the acts he would have faced. Two sorts of conditions will need to be proposed on $\mathcal{S}$ to obtain the representation. On the one hand, there will be general, more or less standard conditions guaranteeing that the agent's choices are consistent and that he is an expected utility maximiser (rather than a non-expected utility maximiser, for example). On the other hand, there will be particular conditions regarding the existence of specific types of situations, which are required for the elicitation technique. The two postulates stating the former conditions will be presented first; then we will turn to the axioms corresponding to the latter conditions.

Basic Postulates According to the definition above, $\mathcal{S}$ may contain two situations with the same states and consequences but different preferences. This represents an agent whose preferences are inconsistent: they may differ on the same set of acts. The theorem below does not apply to such agents. It employs a minimal consistency constraint on the agent's preferences which implies that he is not of this type: namely, that he can only have one order of preference on any set of acts envisaged.

Postulate III.1. For all $\sigma_1, \sigma_2 \in \mathcal{S}$, if $S_{\sigma_1} = S_{\sigma_2}$ and $C_{\sigma_1} = C_{\sigma_2}$, then $\preceq_{\sigma_1} = \preceq_{\sigma_2}$.

This constraint allows the situations to be thought of extensionally; that is, any two distinct situations differ either in their sets of states or in their sets of consequences (or both).

Remark III.1. There is another consistency condition which may be imposed, and which relates to the possibility of "identifying" acts between situations. An act is a function from states to consequences: there is thus no (evident) way of saying that two acts are the same if they belong to situations with different sets of states (they have different domains). By contrast, it does seem that one could identify acts which belong to situations with the same set of states and which give the same consequences (they have the same domain and the same image), although the sets of consequences available in the situations to which they belong differ (they have different ranges). Consider for example $\sigma = (S_\sigma, C_\sigma, \preceq_\sigma)$ and $\sigma' = (S_{\sigma'}, C_{\sigma'}, \preceq_{\sigma'})$, where $S_\sigma = S_{\sigma'}$ and $C_\sigma \subset C_{\sigma'}$: it seems intuitive to identify an act in σ with the act in σ' that gives the same consequences for each state, and thus to demand that the preferences over the acts in σ are the same as the preferences over the corresponding acts in σ'. This is expressed by the following axiom.

Axiom III-A1. For any pair of situations $\sigma_1 = (S_{\sigma_1}, C_{\sigma_1}, \preceq_{\sigma_1})$ and $\sigma_2 = (S_{\sigma_2}, C_{\sigma_2}, \preceq_{\sigma_2})$ with $S_{\sigma_1} = S_{\sigma_2} = S$ and $C_{\sigma_1} \cap C_{\sigma_2} \neq \varnothing$, $\preceq_{\sigma_1} |_{\mathcal{A}_{\sigma_1} \cap \mathcal{A}_{\sigma_2}} = \preceq_{\sigma_2} |_{\mathcal{A}_{\sigma_1} \cap \mathcal{A}_{\sigma_2}}$.

(Standard mathematical notation shall be employed whereby, for an order $\leqslant$ on a set Y and $Y' \subseteq Y$, $\leqslant |_{Y'}$ is the restriction of $\leqslant$ to Y'.)

This axiom shall *not* be necessary for the results in this section, and shall not be assumed to hold unless explicitly stated. However, we shall have reason to discuss it in Sections III.3 and III.4.

The decision-theoretic framework used here is that proposed in Anscombe and Aumann (1963), which shall now be briefly summarised.

For a given situation σ, S_σ is assumed to be finite, and C_σ is assumed to be the set of lotteries over a finite set X_σ – the set of outcomes. Acts – functions from states to lotteries – are thought of as functions from $S_\sigma \times X_\sigma \rightarrow \Re$; that is, the set of acts is $\mathcal{A}_\sigma = \{f : S_\sigma \times X_\sigma \rightarrow \Re | \sum_{x \in X} f(s,x) = 1\}$.[4] Under these assumptions, $\mathcal{A}_\sigma$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}_\sigma$ and $a \in \Re$, $0 < a < 1$, the mixture $af + (1-a)h$ is defined by $(af + (1-a)h)(s,x) = af(s,x) + (1-a)h(s,x)$ (Fishburn, 1970, Ch 13). So the von Neumann-Morgenstern axioms can be stated for the preference orders $\preceq_\sigma$.

Axiom III-A2 (Weak order)**.** (a) For all f, g in $\mathcal{A}_\sigma$, $f \preceq_\sigma g$ or $g \preceq_\sigma f$. (b) For all f, g and h in $\mathcal{A}_\sigma$, if $f \preceq_\sigma g$ and $g \preceq_\sigma h$, then $f \preceq_\sigma h$.

Axiom III-A3 (Independence)**.** For all f, g and h in $\mathcal{A}_\sigma$, and for all $a \in \Re$, $0 < a < 1$, if $f \preceq_\sigma g$ then $af + (1-a)h \preceq_\sigma ag + (1-a)h$.

Axiom III-A4 (Continuity)**.** For all f, g and h in $\mathcal{A}_\sigma$, if $f \preceq_\sigma g$ and $g \preceq_\sigma h$, then there exist $a, b \in (0,1)$ such that $af + (1-a)h \preceq_\sigma g$ and $g \preceq_\sigma bf + (1-b)h$.

Anscombe and Aumann (1963) add the following axiom:

Axiom III-A5 (Monotonicity)**.** For any elements c_1 and c_2 of C_σ, any f in $\mathcal{A}_\sigma$ and any s in S_σ, if $\mathbf{c_1} \preceq_\sigma \mathbf{c_2}$ then $f_s^{c_1} \preceq_\sigma f_s^{c_2}$

where $\mathbf{c}$ is the constant act taking value c (for all $s \in S_\sigma$, $f(s,x) = c(x)$) and f_s^c is identical to f, except on s, where it takes value c (that is, $f_s^c(s,x) = c(x)$ and, for $s' \neq s$, $f_s^c(s',x) = f(s',x)$).

4. In thinking of acts in this way, Anscombe and Aumann's "Reversal of Order" axiom is assumed to hold (Anscombe and Aumann, 1963). As is standard in much work with this framework, this axiom shall be assumed to hold in all situations, throughout this paper. The consequences of rejecting it are explored by Drèze (1987), for example.

Fishburn (1970, p146) has shown that Weak Order, Independence and Continuity imply that there is an additively separable representation of $\preceq_\sigma$ (a representation of form (III.1); see Section III.1), such that the evaluation function is unique up to simple positive affine transformation: for any pair of functions U and U' representing $\preceq_\sigma$, there is a positive real number a and real numbers b_s for each $s \in S_\sigma$ such that $U'(s,x) = aU(s,x) + b_s$ for all $s \in S_\sigma$ and $x \in X_\sigma$. Anscombe and Aumann (1963) show furthermore that adding the Monotonicity axiom yields a decomposition of this function into a unique probability function over the set of states and a state-independent utility function over the set of outcomes which is unique up to positive affine transformation. Indeed, the Monotonicity axiom is so closely connected to state independence that it is often simply called the axiom of state independence (in Hammond (1998), for example).

The current work will follow much of the literature on state dependence of utilities and suppose that the expected utility axioms other than those yielding state independence hold in every situation. This is expressed by the following postulate.

Postulate III.2. For any situation $\sigma \in \mathcal{S}$, $\preceq_\sigma$ satisfies Weak Order, Independence and Continuity.

In the light of what has been said above, this postulate implies that the agent always maximises an additive evaluation function, whether or not this function can be decomposed (in a non-arbitrary manner) into a probability function and a utility function. Behaviourally, an agent satisfying this postulate always acts as if he is an expected utility maximiser (rather than a non-expected utility maximiser), but, from an elicitation point of view, satisfying the postulate is not enough to allow identification of his probabilities and utilities.

The situations where there exists a decomposition of the evaluation function into a probability and state-independent utility function are of special importance.

Definition III.2. σ is called a *simple* situation if and only if $\preceq_\sigma$ also satisfies Monotonicity.

The simple situations are those which permit the application of the Anscombe and Aumann representation theorem. Let us state this explicitly.

Theorem III.1 (Anscombe and Aumann (1963))**.** *For a simple situation σ, there is a unique probability distribution p_σ on S_σ and a utility function u_σ on X_σ which is unique up to positive affine transformation, such that, for $f, g \in \mathcal{A}_\sigma$*

(III.3) $$f \preceq_\sigma g \textit{ iff } \sum_{S_\sigma, X_\sigma} p_\sigma(s) u_\sigma(x) f(s,x) \leqslant \sum_{S_\sigma, X_\sigma} p_\sigma(s) u_\sigma(x) g(s,x)$$

An unsympathetic reading of the literature would have it that *all* situations are simple, so that theorems such as this one can always be applied (Savage (1954, §5.5), for example, seems to assume that his axioms apply in all "small world" decision situations). This cannot assumed here, for the challenge is to elicit probabilities and utilities in non-simple situations (i.e. situations where the state-independence axioms do not apply). A milder assumption will be made: that there are *enough* simple situations to permit elicitation of (unique) probabilities and utilities. The axioms in the representation theorem will cash out this assumption in formal terms. First it is necessary to introduce some preliminary definitions.

Definition III.3. For a set of consequences C, let $\Sigma_C = \{\sigma \in \mathcal{S} | \; C = C_\sigma, \sigma \text{ simple}\}$. This is the set of simple situations with C as set of consequences.

Define $\|\Sigma_C\| = \{u | \; \exists \sigma \in \Sigma_C, p \text{ on } S_\sigma, \text{ s.t. } p, u \text{ represent } \preceq_\sigma\}$, the set of utilities involved in the representation of the preferences (according to (III.3)) in the simple situations. By Theorem III.1 and Definition III.2, $\|\Sigma_C\|$ is non-empty if Σ_C is. The elements of this set will be considered up to positive affine transformation, to avoid unnecessary repetitions.

Similarly, for a set of states S, define $\Xi_S = \{\sigma | \; S = S_\sigma, \sigma \text{ simple}\}$, and $\|\Xi_S\| = \{p | \; \exists \sigma \in \Xi_S, u \text{ on } C_\sigma, \text{ s.t. } p, u \text{ represent } \preceq_\sigma\}$.

The idea is to use the representations of the preferences in simple situations to ascertain the utility in non-simple situations which have the same set of consequences; similarly for probabilities. For this several existence and uniqueness constraints are required on the set of simple situations; these will be the main axioms in the representation theorem. The axioms are presented here; further discussion is left for Section III.3.

Main axioms Let us first deal with the situations sharing the same consequences. Consider the following two axioms.

Axiom III-A6 (Consequence Richness)**.** For all C, Σ_C is non-empty.

Axiom III-A7 (Conative Consistency)**.** For all C, for any $\sigma_1, \sigma_2 \in \Sigma_C$, $\preceq^{const}_{\sigma_1} = \preceq^{const}_{\sigma_2}$, where $\preceq^{const}$ is the restriction of $\preceq$ to the constant acts.

Consequence Richness states that there are enough situations such that, for each set of consequences, an "independent" set of states can be found – independent in the sense that the Anscombe and Aumann's state-independence axiom, and thus their expected utility result, applies. As such, it can be considered to be a structural or technical axiom: its behavioural consequences are negligible, and, due to the existential quantifier, it is difficult to refute. From a technical point of view, this axiom guarantees that $\|\Sigma_C\|$ has at least one element. See Section III.3 for an extended discussion of the plausibility of the axiom.

Conative Consistency is equivalent to the existence of at most one element in $\|\Sigma_C\|$ (Lemma III.A.1 in the Appendix). By contrast to the previous axiom, it does have some behavioural content: it demands that the agent's preferences over constant acts – acts which give the same result on all states – is independent of what the states are (and thus, given Definition III.1 and Postulate III.1, of what the situation is). An agent who prefers \$50 to a 50-50 objective lottery yielding \$100 when in the context of a horse race (states: winners of the race) but prefers the objective lottery to the sure amount in the context of investment on the markets (states: stock price tomorrow) violates this axiom. Certainly, such preferences do not seem particularly consistent: he prefers receiving the sure amount, no matter who wins the horse race, to the playing objective lottery, no matter who wins the horse race, but he prefers the lottery, no matter what the stock price is tomorrow, to the getting the sure amount, no matter what the stock price is tomorrow.

One proceeds in a similar way for states and probabilities; once again, there are two axioms.

Axiom III-A8 (State Richness)**.** For all S, Ξ_S is non-empty.

Axiom III-A9 (Doxastic Consistency)**.** For all S, for any $\sigma_1, \sigma_2 \in \Xi_S$, for any $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ with $\mathbf{a} \prec_{\sigma_1} \mathbf{b}$ and $\mathbf{c} \prec_{\sigma_2} \mathbf{d}$, let $\tau : L(\{a, b\})^S \rightarrow L(\{c, d\})^S$ be the bijection between the set of acts in σ_1 with values in the set of lotteries on $\{a, b\}$ and the set of acts in σ_2 taking values in the set of lotteries on $\{c, d\}$, defined by $\tau(f)(s, c) = f(s, a)$ and $\tau(f)(s, d) = f(s, b)$ for all $s \in S$.[5] Then, for all $f, g \in L(\{a, b\})^S$, $f \preceq_{\sigma_1} g$ iff $\tau(f) \preceq_{\sigma_2} \tau(g)$.

State Richness is the equivalent for states of Consequence Richness (III-A6): it ensures that $\|\Xi_S\|$ has at least one element. As for III-A6, this is a largely structural assumption.

5. $L(X)$ is the set of lotteries on the set X; $L(X)^S$ is the set of functions from S into $L(X)$ – that is, the set of acts taking values in these lotteries. Recall that the acts are considered as functions from pairs of states and outcomes to the real numbers.

Doxastic Consistency is the equivalent of Conative Consistency (III-A7): it is true if and only if $\|\Xi_S\|$ has at most one element (Lemma III.A.2). As for Conative Consistency, it does have behavioural content: more or less, it demands that, to compare probabilities of events (sets of states) using acts which differ in their consequences depending on whether the event holds, it does not matter if one uses consequences in C_{σ_1} or C_{σ_2} – the answer will be the same. Consider two situations: one where the agent is betting on a horse race (states: results of the race; consequences: monetary prizes) and another where he is advising his friend how to bet on the race (states: results of the race; consequences: monetary prizes for the friend). Suppose furthermore that the agent prefers a sure $\$100$ to a sure $\$0$ in the first situation, and that he prefers that his friend wins $\$100$ for sure to his friend winning nothing for sure in the second situation. Then, according to Doxastic Consistency, he should prefer a bet of $\$100$ on horse α over a bet on horse β if and only if he prefers advising his friend to take the bet on α to advising his friend to take the bet on β. Indeed, an agent who did not have such preferences would seem rather inconsistent.

Given the preceding remarks, there is a resemblance between Doxastic Consistency (III-A9) and Savage's P4, which states that the preferences over bets on events are independent of the consequences of the bets. But P4 holds *within* simple situations, whereas Doxastic Consistency holds *between* different simple situations. So, if Doxastic Consistency holds between situations having different sets of consequences, one might expect P4 to hold in the situation having the union of these sets as consequences. This is essentially what is expressed by the following axiom, which implies Doxastic Consistency, in the presence of III-A1 (Proposition III.1 in the Appendix).

Axiom III-A10. For any S, if $\sigma_1, \sigma_2 \in \Xi_S$, then there is a situation $\sigma_{12} \in \Xi_S$ with $C_{\sigma_{12}} = C_{\sigma_1} \sqcup C_{\sigma_2}$, where $C_1 \sqcup C_2$ is the set of lotteries on $X_1 \cup X_2$.

This axiom is not required for the result stated below; it shall however prove relevant in Section III.4.

Null events One final axiom is required to deal with the possibility of null events. In the state-independent representation, null events are generally those which are allocated probability 0: it follows that the order is indifferent between any pair of acts which differ only on such events. For the probabilities elicited in simple situations to be valid in non-simple situations, the preference order in these non-simple situations must show the same sort of indifference. Thus the following axiom is posed as a consistency constraint. Recall the classic definition of a null event (Savage, 1954; Karni et al., 1983): an event A in a situation σ is *null* iff, for any pair of acts $f, g \in \mathcal{A}_\sigma$ such that $f(s) = g(s)$ for $s \notin A$, $f \sim_\sigma g$. (Null states are those whose singletons are null events.) The axiom is as follows.

Axiom III-A11 (Null Consistency)**.** For any situation σ and any event $A \subseteq S_\sigma$, if A is null in every $\sigma' \in \Xi_{S_\sigma}$, then A is null in σ.

III.2.2 Theorem

The postulates and axioms give the following representation theorem.

Theorem III.2. *Assume Postulates III.1 and III.2. Moreover, let Consequence Richness, Conative Consistency, State Richness, Doxastic Consistency and Null Consistency hold. Then, for any situation $\sigma \in \mathcal{S}$, there exists a probability distribution p on S_σ, a utility function u on X_σ, and a function $\gamma : S_\sigma \times X_\sigma \to \Re$ such that,*

— *for all $f, g \in \mathcal{A}_\sigma$*

$$f \preceq_\sigma g \text{ iff } \sum_{\substack{s \in S_\sigma \\ x \in X_\sigma}} p(s).\gamma(s,x).u(x).f(s,x) \leqslant \sum_{\substack{s \in S_\sigma \\ x \in X_\sigma}} p(s).\gamma(s,x).u(x).g(s,x) \tag{III.4}$$

— *for each $\sigma_1 \in \Sigma_{C_\sigma}$, there exists a probability p_1 such that, for all $f_1, g_1 \in \mathcal{A}_{\sigma_1}$*

$$f_1 \preceq_{\sigma_1} g_1 \text{ iff } \sum_{\substack{s \in S_{\sigma_1} \\ x \in X_{\sigma_1}}} p_1(s).u(x).f_1(s,x) \leqslant \sum_{\substack{s \in S_{\sigma_1} \\ x \in X_{\sigma_1}}} p_1(s).u(x).g_1(s,x) \tag{III.5}$$

— *for each $\sigma_2 \in \Xi_{S_\sigma}$, there exists a state-independent utility u_2 such that, for all $f_2, g_2 \in \mathcal{A}_{\sigma_2}$*

$$f_2 \preceq_{\sigma_2} g_2 \text{ iff } \sum_{\substack{s \in S_{\sigma_2} \\ x \in X_{\sigma_2}}} p(s).u_2(x).f_2(s,x) \leqslant \sum_{\substack{s \in S_{\sigma_2} \\ x \in X_{\sigma_2}}} p(s).u_2(x).g_2(s,x) \tag{III.6}$$

Furthermore, if p', u', γ' is another representation satisfying (III.4–III.6), then there exist positive real numbers a and c, and real numbers b and d_s for each $s \in S_\sigma$, such that $p'(s) = p(s)$, $u'(x) = a.u(x) + b$ and $\gamma'(s,x) = c.\gamma(s,x) - \frac{b.c}{u'(x)}.\gamma(s,x) + \frac{d_s}{u'(x)}$, for all $s \in S_\sigma, x \in X_\sigma$, and $d_s = 0$ if s is null.

Remark III.2. The probability p and utility u have the same uniqueness properties as in typical state-independent utility theorems (for example Theorem III.1); to this extent, they can be said to be elicited in Theorem III.2. γ has new degrees of freedom which arise largely from the fact that, whereas state-independent utilities are unique up to positive affine transformation, state-dependent utilities are unique up to simple positive affine transformation (see for example Karni et al. (1983), who call these transformations "cardinal unit comparable transformations"). Indeed, the representation (III.4) naturally yields a state-dependent utility function as the product of γ and u (see Section III.3.1), and it is easy to see that this function is unique up to simple positive affine transformation, just as for the state-dependent utilities elicited by Karni et al. (1983). That is, for another triple p', γ', u' satisfying the conditions of Theorem III.2, there is a positive real number a' and real numbers b'_s for each state $s \in S_\sigma$ such that $\gamma'(s, x).u'(x) = a'\gamma(s, x).u(x) + b'_s$.

Remark III.3. The theorem does not assume III-A1, and so allows differences between the preferences on acts taking the same set of states to the same consequences, but which are elements of different situations (notably, situations having the same sets of states and different but overlapping sets of consequences). In such cases, the representations of the preferences in the different situations may differ in γ and u.

On the other hand, if III-A1 holds, the preferences agree (on common acts) between such situations; it follows that u takes the same values in situations with at least two outcomes in common and γ takes the same values in all situations which share the same set of states and which have at least two outcomes in common (up to the transformations given in Theorem III.2).

III.3 Discussion

III.3.1 The situation-dependent factor

The factor γ is a function of states and consequences. Since different situations necessarily have different states and consequences, they will necessarily involve different γ's. Thus γ is properly thought of as a situation-dependent factor: particular to the decision situation under consideration, rather than applicable in different decision situations, as, say, utilities and probabilities are. Put succinctly, γ captures the contextual factors of the decision situation.

It should be noted that a single decision problem can sometimes be rewritten in several ways. The discussion in Section III.1 provides an example: the problem the husband faces can be formulated either as a choice between acts yielding one of two consequences (Figure III.1) or as a choice among acts yielding one of four consequences (Figure III.2). These formulations of the decision problem are two different situations, in the sense of Definition III.1, with different sets of consequences; therefore, the representations will involve different factors γ.

Consider, for example, under the 4-consequence formulation (Figure III.2), 50-50 (subjective) probabilities, a utility of 150 for $\$100$ with the wife alive, of 50 for $\$100$ with the wife dead, of 75 for $\$0$ with the wife alive and of 25 for $\$0$ with the wife dead, with γ taking value 1 when the state and consequence are compatible, and 0 when they are not. This is a typical representation of the husband's preferences, which explains the preference for the bet on success (bet A in Figure III.2) over failure (bet B) by a larger utility of $\$100$ with the wife alive, and the preference of the bet on success (bet A) over a bet on resuscitation (bet E) by a low value of γ which represents the fact that the latter bet is unbelievable. However, it is equally possible to represent the husband's preferences when the 2-consequence formulation is used (Figure III.1), just this will require a different utility function (because the consequences are different) and a different situation-dependent factor γ. For example, 50-50 probabilities, a utility of 100 for $\$100$ and 50 for $\$0$, and a γ with $\gamma(\text{success}, \$100) = 1.5$, $\gamma(\text{failure}, \$100) = 0.5$, $\gamma(\text{success}, \$0) = 1.5$, $\gamma(\text{failure}, \$0) = 0.5$ represents "the same" preferences as described above in the 4-consequence formulation.[6] Notably, they account for the preference for the bet on success, whilst retaining the 50-50 beliefs about the result of the operation. The use of the 2-consequence formulation or the 4-consequence formulation is largely a decision for the theorist; as this example indicates, the proposed representation (III.2) applies whatever choice is made.

As for all theorems treating decision under uncertainty, Theorem III.2 elicits the probabilities and utilities; moreover, it also elicits the situation-dependent factor γ. The latter is thus to be thought of as "subjective" rather than "objectively" given. It represents aspects of the way that the agent thinks of the decision problem and the contextual factors involved,

6. Naturally, any claim that the preferences are the "same" has to be qualified by the fact that, because of the different sets of consequences, the objects of preference – the acts – differ between the two formulations. For further discussion of the relationships between the situation-dependent factors and utilities in such different but related situations, see Section III.3.2.

rather than the contextual factors "really" involved. For example, if the decision maker did not realise that resuscitation was impossible, the γ described above (for the 4-consequence formulation of the example) would not feature in an accurate description of his preferences (for he would be indifferent between bet A and bet E in Figure III.2), even if it is an objectively accurate representation of the relationship between states and consequences. The topic of objective situation-dependent factors is beyond the scope of this paper.

A situation-dependent factor of this sort, and a representation of the form (III.2) proves useful in the analysis of the phenomenon of adaptive preferences. In particular, it allows a distinction between the agent's "absolute" utility u – which is independent of the situation in which he finds himself – and his "situation-relative" utility or "utility in practice". The latter is the state-dependent utility function obtained as the product of the factor γ and the absolute utility u. See Hill (2009) for an extended discussion of this distinction and the importance for the problem of adaptive preferences. [7]

The role of γ as characterising the interdependence between states and consequences, or the difference between the "absolute" and "situation-relative" utilities may prove useful in other areas of economic analysis. It is known that the agent's attitude to risk may differ depending on the state realised: his attitude to risk is different in the case of illness as opposed to health, or in the case of life as opposed to death (Karni (1983), Drèze (1987, Ch 8)). Under state-dependent utility representations of his preferences, this difference is built into the utility function: the function has different properties (form, curvature, and so on) on each state. Under the proposed representation, by contrast, it is separated from his general attitude to risk. The "absolute" utility represents the agent's general attitude to risk (in the standard way), whereas the situation-dependent factor γ captures the aspects which are specific to the interdependence between states and consequences involved in the particular decision situation.

Consider the example of health insurance, to which state-dependent utilities are often applied: the states are states of health and the consequences are monetary outcomes. γ represents the differences in utility and risk attitude with respect to change in state of health, whereas the utility function measures his utility for money (and the risk attitude associated) and applies even beyond decisions regarding health insurance (to investment decisions,

7. Another possible interpretation is to think of u as the "intrinsic" utility and the product with the situation-dependent factor as the "experienced" utility. The author wishes to thank a referee for pointing out this interpretation, and the relation to the distinction drawn by Kahneman et al. (1997).

career decisions, and so on). For an agent who is more risk averse in case of ill health than good health, this fact would be represented by the fact that γ is more concave when the state is ill health than when it is good health. Future research would investigate whether tools for the analysis of risk attitude in cases where there is interdependence between states and consequences can be developed employing the situation-dependent factor proposed here. One would expect to obtain techniques for doing comparative statics on the situation-dependent factor: one could thus compare the different attitudes to health insurance of individuals with the same beliefs about their future health and the same general utility for money. [8]

It should be clear from the discussion so far that, beyond the observation that γ is a contextual or situation-dependent factor, there are several ways to give it a more concrete interpretation. A full discussion of the range of interpretations is beyond the scope of this paper; we shall just present one further interpretation which is perhaps applicable in some if not all cases, and which has already been suggested in the discussion of adaptive preferences (Hill, 2009).

Consider once again the 4-consequence formulation of Aumann's example. As noted above, the situation-dependent factor which corresponds to the described preferences typically takes lower values on pairs of states and consequences which are inconsistent (failed operation and $\$100$ with wife alive, for example) than on pairs which are consistent. A natural interpretation of this is in terms of reliability. The act purporting to take the state where the operation fails to the consequence where he wins \$100 and his wife is healthy is not possible. If offered such an act (or such a bet), the agent would not trust the bookey. The act is *unreliable* (or, more precisely, the part of the act purporting to send this state to this consequence in unreliable). The expected utility should take this into account, and it is precisely the situation-dependent factor γ that does this. An interpretation of the situation-dependent factor which seems plausible in at least some cases is as a measure of the reliability of acts purporting to take particular states to particular consequences. It should be noted finally that under this interpretation it is most natural to expect γ to have maximum and minimum values (utterly reliable and utterly unreliable). It is not too difficult to show, by scaling the

8. It may be perhaps be possible to use differences in the situation-dependent factor to account for the observation that agents' risk aversion depends on the circumstance or domain (MacCrimmon and Wehrung, 1990; Weber et al., 2002); if so, the notion might also prove useful in studies of such differences in risk aversion.

utilities so that they are positive, that Theorem III.2 can always yield a representation with $\gamma \in [0, 1]$. This form of the theorem is not given here, for such a range of γ may not be natural for other interpretations. Indeed, it is worth re-emphasising that this interpretation is one of many: to the extent that γ represents the effect of contextual or situation-dependent aspects on the agent's preferences, there may be as many interpretations are there are sorts of contextual or situation-dependent factors.

III.3.2 The elicitation technique

The elicitation technique is based on a simple idea: instead of trying to elicit attitudes in cases which are complicated (for example, where there is interdependence between states and consequences), use cases which are straightforward (where there is no such interdependence). Such a technique invokes situations other than the one in question; it thus involves not a single situation, but a set of situations, $\mathcal{S}$.

The situations in $\mathcal{S}$ are the decision problems the agent could have been facing at the particular moment in question (given by the sets of states and consequences) and the preferences he would have had were he faced with these problems (given by the preference relation). One of the situations is actual: that is the situation which he is really facing at a given moment. The others are hypothetical: he could have faced them, and would have had the preferences specified if he had, but he is not actually facing them (see below for further discussion of the notion of hypothetical choice involved here). Some of these situations are simple, insofar as they do not involve interdependence between states and consequences: they are used to elicit the agent's probabilities and utilities. The method of elicitation relies on a basic assumption: that the agent's attitudes (probabilities and utilities) have a certain degree of stability or constancy. For it to make sense to elicit the agent's probabilities and utilities in the actual situation by looking in other situations, it must be assumed that the same probabilities and utilities are involved in different situations.

This assumption is not only natural, but a necessity for the task of elicitation or measurement to make any sense. First of all, one would expect rational agents to have probabilities and utilities which are stable enough not to be excessively dependent on the precise decision problem with which they are faced; and many real agents' attitudes do have this property. Secondly, if one does not suppose such constancy, one cannot use the measurements of an agent's probabilities and utilities in any situation other than that in which it was measured. For example, probabilities and utilities measured in an experiment could not be used in models of agents' behaviour in the market. There is a risk that this problem does indeed affect some experimental work (Harrison et al., 2007); nevertheless, for there to be any interest in measuring attitudes, an assumption such as constancy seems to be necessary.

Furthermore, many elicitation techniques do in fact seem to make implicit use of an assumption like the constancy assumption. Often different sorts of questions are used to elicit different components of the representation (probabilities, utilities, decision weights and so on); these different questions correspond, in the terminology of this section, to different situations (see Section III.4 for a construal of these situations as subsituations of a single fixed situation). One must thus assume that the probability function elicited in one situation is the same as the one involved in the other situation. Naturally, were an experimental setup to be proposed based on Theorem III.2, it would use exactly this method. First the agent's probabilities would be elicited using standard techniques but in situations which do not necessarily involve the consequences in the decision problem of interest. Then the same would be done for utilities. Finally, by eliciting his evaluation function in the situation of interest, the situation-dependent factor would be derived.

Beyond being essential to the elicitation technique employed in the theorem, the constancy assumption does have some behavioural content. Indeed, the Conative and Doxastic Consistency axioms (III-A7 and III-A9) – which constitute, along with the requirement of basic consistency of preferences (Postulate III.1) and the requirement that the agent's preferences have an additive representation (Postulate III.2), the main behavioural content of the theorem – follow from the constancy assumption. If the probabilities and utilities are the same in all situations, then in those situations where they can be elicited, one obtains the same answer. The axioms imply that this is the case: that the utility and probability elicited in any simple situations are always the same, for a given set of consequences and states.

The elicitation technique rests upon another assumption, regarding the richness of the set of situations $\mathcal{S}$. To elicit the agent's probabilities in a given non-simple situation, the technique looks in other situations, which have the same states, but different and independent consequences, so that elicitation of probabilities is simple. Hence it needs to be assumed that such situations exist. Similarly, to elicit utilities, it needs to be assumed that simple situations with the same consequences but perhaps different states exist. This is basically the content of the Consequence and State Richness axioms (III-A6 and III-A8). As noted in Section III.2.1, these can be thought of as structural or technical; nevertheless, they are not unreasonable. To take the example discussed in Section III.1 (in the 2-consequence version), consider the decision situation where the states of the world are results of a horse race and the consequences are $\$100$ or $\$0$. In this situation the consequences are most likely independent of the states (assuming the husband does not know any of the horse-owners or jockeys), so that the state-independence axioms apply and the utility of the money can be elicited using standard techniques. This situation has the same consequences but different states from the situation where the agent bets on the outcome of the operation, so that the former situation can be used to elicit the utilities in the latter one. Hence Consequence Richness (III-A6) is satisfied in this case. Similar examples can be given for states and the axiom State Richness (III-A8). All that is required are situations with the same states (success or failure of the operation) but consequences which are independent, so that the state-independent axioms and results apply. Examples include situations were the consequences are gains and losses for a friend, and the husband has a choice about the advice he gives (see also Karni (1996, p256)), or situations where the consequences of the husband's bet will be "ethical" issues, such as differing numbers of deaths in Myanmar.[9] In cases such as these, the utility of the consequence does not depend on the outcome of the operation (the state of the world), and so the state-independence axioms apply; they show that, for the example in Section III.1, State Richness is satisfied.

The same points hold for the 4-consequence formulation of the decision problem. The utility function elicited in the 4-consequence formulation will be different from that elicited in the 2-consequence formulation because the sets of consequences on which these functions are defined differ. Equally, the situations used to elicit them will differ: in one case, bets on a horse race yielding monetary prizes will be involved, in the other case, bets yielding combinations of monetary prizes with a deceased or healthy wife will be used. Naturally, one might sometimes expect there to be relations between the utility functions in these different situations, just as one might expect there to be relations between the situation-dependent factors, which also differ between the situations because of their different consequences (Section III.3.1). Candidate relations would ideally come in the form of equations which connect the utilities, probabilities and state-dependent factor in one sit-

9. Thanks to Itzhak Gilboa for suggesting this example.

uation to those in the other situation, and which correspond to behavioural conditions on the agent's preferences – notably relations between his preferences in the two situations. A simple example of such a relation has already been mentioned. Axiom III-A1 presents a condition on the preferences of the agent in situations with different but overlapping sets of consequences; as noted in Remark III.3, it ensures that the utilities on common consequences, and the values of the situation-dependent factors on common states–consequence pairs are the same in the different situations. A direction for future development would consist in finding and assessing such conditions on preferences in different situations, and the corresponding consequences for the relation between utilities, probabilities and state-dependent factors in these situations.

Such a project requires a clear idea of the interpretation of the situations, and of the differences between them. Consider once again Aumann's example and the elicitation of utilities by reference to situations where the agent is betting on a horse race. In the situation with 4 consequences, the outcome of the wife's operation is explicit in the consequences, whereas in the situation with 2 consequences, it is not – if it plays a role in the evaluation of the consequences, this role is at most implicit. The explicitness or implicitness is an important aspect of the relationship between the two situations, and may be interpreted in at least two ways.

Firstly, it may be understood as a difference in the way the problem is posed to the agent: he has his wife's health in mind at all times, but has to make choices where the health of his wife is not explicitly at issue. In such cases, one might expect a principled relationship between his utilities in the 2-consequence situation and those in the 4-consequence situation. For example, he might be expected to consider each consequence in the 2-consequence situation (say, $\$100$) as a bet over pairs of consequences in the 4-consequence situation ($\$100$ with the wife alive and $\$100$ with the wife dead), so that his utility for the 2-consequence outcome is the expected utility of $\$100$ over his subjective probabilities about the outcome of the operation (Savage, 1954, §5.5).

Alternatively, the implicitness and explicitness could be interpreted as properties of the way the agent himself sees the problem. The 4-consequence formulation corresponds to the case where he is aware of the effect of the operation's success or failure on his utility for money; the 2-consequence formulation represents the case where he is not aware of the issue of the operation. In this case, it is possible that his utilities, probabilities and situation-dependent factor are related in the sort of principled way just described: although he is unaware, he evaluates the options in the same way as if he were aware. On the other hand, it more likely under this interpretation that the relationships between the utilities, probabilities and state-dependent factors in the 2- and 4-consequence situations are more complicated. [10]

A further point should be made about this last interpretation. There are many things about which the agent is not necessarily certain and which could affect his utility for money: the health of himself and his family, his plans and ambitions, his current and future state of wealth, the relation to the state of the economy, and so on. It cannot be asked of him that he be aware of all these issues when he chooses acts yielding monetary consequences; at most, they play an implicit role. So one cannot reject the 2-consequence formulation of the husband's decision, the situation required for elicitation of the utility in this formulation, and the utility elicited, on the basis that a relevant factor – the outcome of the operation – was "left out"; for there are more relevant factors than could ever feasibly be explicitly "brought in". The elicitation of utility described yields a utility function in which these factors are, so to speak, implicitly present.

Let us remark finally that this is not the first time that situations other than the one of immediate interest have been used in representation theorems. Examples include Karni (2007, 2011) and Karni et al. (1983); Karni (1985). In the former cases, reference is made to the situation the agent finds himself in after having made a Bayesian update; in the latter cases, the situations involved are hypothetical. To this extent, the latter cases resemble the proposal here, where, as stated above, the natural interpretation of the other situations involved is as hypothetical. Indeed, some of the arguments presented by Karni and Mongin (2000) in favour of the use of hypothetical data in elicitation carry over to the current result; see Hill (2009) for other arguments which are particular to sort of data involved here.

However, beyond this surface resemblance, there is an important difference between the approaches just cited and the one taken here. There, the other situations which are used involve the *same decision problem* (states and consequences) but *different* (hypothetical or updated) *beliefs*, whereas in this paper, we have (to the extent that the states and consequences are shared) the *same beliefs and utilities* but *different decision problems*. This

10. Awareness has recently been recognised as important in decision and game theory (Dekel et al., 2001; Heifetz et al., 2006). So far, no accepted theory exists of the relationship between the agent's attitudes when he is aware and those when he is not. See Hill (2010) for an approach.

difference is crucial, for it means that the papers mentioned above involve elicitation of attitudes with reference to preferences which are incompatible with the preferences the agent actually has, just as for Savage's strategy in the example discussed in Section III.1. Here, by contrast, there is no incompatibility between the preferences in the situations involved in the elicitation: as the previous discussion of the constancy assumption makes clear, this point is at the core of the present approach. The sense in which there is something hypothetical at issue is thus milder here: whereas previous approaches have considered what would happen with a different decision maker (with hypothetical beliefs) faced with the same decision problem, we are considering what would happen if the same decision maker were faced with a different decision problem. As suggested above, the second sense of hypothetical is easier to overcome in practice: it suffices to get the decision maker to choose in a decision problem in which it is easier to elicit his attitudes.

A further consequence of this difference is that, since there is no incompatibility among his preferences, they can all in principle be considered as "fragments" of a (consistent) set of general preferences. From the point of view of the general preferences, the decision situations we have defined are just fragments of a large decision problem. Since this decision problem is actual, the distinction between the hypothetical and actual decision situations disappears – they are just subsituations of one, large, actual situation. One thus obtains a version of Theorem III.2 which makes no reference at all to hypothetical situations: a relief for those who are not convinced by the preceding arguments in favour of the mild notion of hypothetical situation employed here. In the final section of the paper, we detail this setup and prove a simple version of this result.

III.4 Non-situational version

III.4.1 Preliminaries and Axioms

The use of situations other than the one in which the decision maker is choosing may be shunned. Fortunately, they are not strictly necessary for the result given above: instead of thinking of the situations used in the result as different situations, one could consider the sets of states, consequences and acts of each situation to be appropriate *subsets* of the elements (events, consequences, acts) of some large, fixed situation in which the decision maker is making his choice. Situations would thus be considered to be like the "microcosms" of Savage (1954, §5.5), so the relation of their "small world" states and consequences to any fixed "grand world" states and consequences can be understood. Whereas the bets on the success and failure of the operation and the bets on the horse race were considered as acts belonging to different situations in the previous sections, here they are considered as being part of one grand world situation, whose states specify both the outcome of the operation and the results of the race. In this way, it is possible to formulate equivalents to Theorem III.2 in terms of one "grand world" situation. In this section, a simple theorem of this sort is stated. It gives a flavour of how the technique introduced in the previous sections continues to apply in a single, grand world situation.

Let S designate the set of states of the grand world situation, X the set of outcomes, C the set of lotteries on X, and $\preceq$ the preference order on the set of acts $\mathcal{A}$ (functions from S to C). S and X are assumed to be finite.

To obtain a representation of the preferences and elicit unique probabilities and utilities, the idea presented and discussed in Sections III.2 and III.3 will be employed: elicit the probabilities and utilities in related, simple, situations. By contrast to the previous sections, the situations used here will be "small world" situations with respect to the fixed grand world situation, or if you prefer *subsituations* of the grand world situation. That is, they will be situations of the form $(S_\sigma, C_\sigma, \preceq_\sigma)$ where S_σ is a set of events which partition S, C_σ is the set of lotteries on X_σ for some $X_\sigma \subseteq X$ and $\preceq_\sigma$ is the restriction of $\preceq$ to acts which are constant on the elements of S_σ and which take values in C_σ. A small world state is a grand world event; the same notation will be used when it is treated as the former as when it is thought of as the latter. An act f in the (small world) situation $(S_\sigma, C_\sigma, \preceq_\sigma)$ is considered as a (grand world) act $\hat{f}$, where, for each $s \in S$, $\hat{f}(s, x) = f(s_\sigma, x)$, for s_σ the element of S_σ such that $s \in s_\sigma$.[11]

Axioms Beyond the axioms introduced in Section III.2, the following axioms will be required.

Axiom III-A12. There exists a set of sets of events on S, $\{S_i\}_{i \in I}$, such that,

(i) for each $s \in S$, there is a set $\{s_i^s\}$, with one element s_i^s from each S_i, such that $\bigcap_{i \in I} s_i^s = \{s\}$;

11. This interpretation differs from Savage's interpretation of microcosms (Savage, 1954, §5.5); thus the difference in the interpretation of small consequences, which are not grand world acts, as in Savage.

(ii) for each S_i, there exists $X' \subseteq X$ such that $(S_i, C', \preceq')$ is simple (where, as above, C' is the set of lotteries on X');

(iii) for each $s \in S$, the set $\{s_i^s\}$ is minimal: there is no proper subset $A \subset \{s_i^s\}$ such that $\bigcap_{s' \in A} s' = \{s\}$.

Axiom III-A13. For each set of events S' which partition S, if there exists $X' \subseteq X$ such that $(S', C', \preceq')$ is simple, then there is a unique such X' which is maximal under inclusion – that is, such that there is no $X'' \supset X'$ with $(S_i, C'', \preceq'')$ simple.

For each S_i from III-A12, the set of outcomes described by III-A13 is denoted by X_i, the corresponding set of consequences by C_i, and the preference order by $\preceq_i$.

Axiom III-A14. For any simple small world situation $(S', C', \preceq')$, if $A \subseteq S'$ is a null event in $(S', C', \preceq')$, then $\bigcup_{s_j \in A} s_j$ is a null event in $(S, C, \preceq)$.

Discussion Just as for the result in Section III.2, two sorts of conditions are required. On the one hand, there are standard conditions guaranteeing that the agent is a consistent expected utility maximiser. As in the previous sections, the axiom for state-independence (Monotonicity) is not assumed to hold in the grand world situation, but the von Neumann-Morgenstern axioms (Weak Order, Independence and Continuity) are (see Theorem III.3). This assumption is logically stronger than the equivalent in the situational setup (Postulate III.2): the former implies the latter, because it implies that the axioms hold in individual situations; but the former is not implied by the latter because it involves inter-situational comparisons. Note furthermore, that, since the preference orders in the small world situations are restrictions of a grand world preference order, they agree on the acts they have in common, so Postulate III.1 and the stronger III-A1 hold automatically.

The other conditions involved in Theorem III.2 guarantee that the set of situations is suitable for elicitation of probabilities and utilities. In the theorem presented below, there are no assumptions or axioms regarding consequences and utilities. This is because the preference relations in the small world situations are derived from the preference relation in the grand world situation, so that they agree on constant acts, and the restriction to constant acts yields, by the von Neumann-Morgenstern theorem, a utility function which applies in all (small world and grand world) situations. So the use of small world situations is not required to elicit utilities; by contrast, they are required for probabilities, and here supplementary conditions have to be imposed.

Axiom III-A12 is the equivalent of State Richness (III-A8) in the grand world situation framework, and, as for State Richness, it is a largely structural or technical axiom. Recall that State Richness guarantees the existence of simple situations with the same set of states as the situation of interest; the probabilities elicited in the former situations is used to deduce the probability in the latter. However, in the grand world version, it is not elicitations in situations with the same set of states which are used to obtain the probabilities in the grand world situation, but rather elicitations in subsituations; thus one needs the existence of a set of simple small world situations which is rich enough that, by eliciting the probabilities in the small world situations, one can deduce the grand world probabilities. This is the role of III-A12. Clause (ii) implies that probabilities can be elicited on each of the sets S_i, which are the sets of states of simple small world situations. Clause (i) implies that these sets of small world states "cover" the set of grand world states. Formally, this clause implies that the Boolean algebra of events of the grand world situation is the smallest Boolean algebra containing the Boolean algebras generated by the $\{S_i\}_{i \in I}$. In other words, every grand world event can be obtained by taking appropriate intersections and unions of events in the small world situations. If this were not the case, it would never be possible to ascertain the probabilities of some of grand world events using just the probabilities elicited in these small world situations.

Clause (iii) translates a simplifying assumption which is made in this section: namely, that one can work with small world situations whose states are stochastically independent. This assumption is at work in the proof of the theorem, although, as is generally the case for central assumptions in representation theorems, it cannot be stated explicitly: after all, the most natural way to express stochastic independence is in terms of the probability function, and such a function is exactly what is to be elicited. Nevertheless, this assumption has some formal consequences, and clause (iii) of III-A12 is one of them. If the states are stochastically independent, then they are logically independent (for each state $s_j \in S_i$, $s'_j \in S_{i'}$, the intersection $s_j \cap s'_j$ is non-empty), and clause (iii) guarantees this. Technically, it implies that the set of states of the grand world situation is the Cartesian product of the sets S_i.

The supplementary assumption that there exists not only a set of simple small world situations whose states cover the grand world state space but that there exists such a set of situations whose states are stochastically independent is made here largely for convenience. First of all, versions of the result stated below continue to hold *mutatis mutandis* when this assumption is weakened appropriately. Furthermore, given that the aim here is merely to indicate that the results in the previous sections can be reformulated in the context of a single situation, rather than enter into complicated technical details, the simplicity permitted by the assumption justifies its use. Finally, the assumption is quite plausible in many contexts. Naturally, one can always invent a grand world situation (set of states, set of consequences and preference on acts) where the assumption of the existence of simple small world subsituations with stochastically independent states will not hold (and in particular, where III-A12 is violated). However, if, as is more often the case, one knows some of the acts, consequences and events involved in a decision problem and one wishes to elicit attitudes, then it is usually possible to find a model of the problem as a single grand world situation which has the required property. Note, for example, that most of the simple situations considered in the previous sections have stochastically independent states: the state of health of the wife is stochastically independent of the result of the horse race, for example. The grand world situation generated by these simple situations has a set of states which is covered by stochastically independent sets of events.[12] Indeed, what was done in previous sections can be thought of as operating in this, single, grand world situation. For these cases, which concern us here insofar as they show that the interpretation in terms of hypothetical situations is unnecessary, the assumption of stochastic independence is not implausible.

Axiom III-A13 does the job of Doxastic Consistency (III-A9); in fact, given the framework, it is equivalent to the stronger III-A10 (Proposition III.2). The comments made in Section III.2 regarding Doxastic Consistency largely hold for III-A13. In particular, it has behavioural content, demanding that the agent have a certain consistency: the probabilities which can be deduced from his preferences in simple situations do not depend on the consequences and thus the simple situation used.

Finally, III-A14 is the equivalent in the current framework of Null Consistency (III-A11). It just says that, if A as an event in the simple situation $(S', C', \preceq')$ is null, then it is null as an event in $(S, C, \preceq)$.

12. For details on how to generate a grand world state space from small world ones, see Hill (2008).

III.4.2 Theorem

One can thus obtain an (essentially) unique representation of the preference relation of the form (III.2) which agrees with the probabilities and utilities elicited in the simple small world situations.

Theorem III.3. *Let S be a set of states, X a set of outcomes, C the set of lotteries over X and $\preceq$ an order on $\mathcal{A}$. Suppose that III-A2-III-A4 hold on $\preceq$. Suppose furthermore that III-A12 holds with set $\{S_i\}_{i\in I}$ and that III-A13-III-A14 hold. Then, there exists a probability distribution p on S, a utility function u on X, and a function $\gamma : S \times X \to \Re$ such that,*

— *for all $f, g \in \mathcal{A}$*

(III.7)
$$f \preceq g \text{ iff } \sum_{s\in S, x\in X} p(s).\gamma(s,x).u(x).f(s,x) \leqslant \sum_{s\in S, x\in X} p(s).\gamma(s,x).u(x).g(s,x)$$

— *There exist positive α and real β such that, for each $x \in X$*

(III.8)
$$\sum_{s\in S} p(s).\gamma(s,x).u(x) = \alpha u(x) + \beta$$

— *For each $i \in I$, there exists positive α_i and real β_i such that, for each $s_j \in S_i$, for each $x \in X_i$,*

(III.9)
$$\sum_{s\in s_j} p(s).\gamma(s,x).u(x) = \alpha_i p(s_j).u(x) + \beta_i p(s_j)$$

Furthermore, if p', u', γ' is another representation satisfying (III.7–III.9), then there exist positive real numbers a and c, and real numbers b and d_s for each $s \in S$, such that $p'(s) = p(s)$, $u'(x) = a.u(x) + b$ and $\gamma'(s,x) = c.\gamma(s,x) - \frac{b.c}{u'(x)}.\gamma(s,x) + \frac{d_s}{u'(x)}$, for all $s \in S, x \in X$, and $d_s = 0$ if s is null.

Remark III.4. (III.8) corresponds to (III.5) of Theorem III.2 insofar as it states that the preferences on constant acts is represented by u (and that this representation is unique up to positive affine transformation).

(III.9) corresponds to (III.6) of Theorem III.2 insofar as it reflects the fact that the preference order in situations $(S_i, C_i, \preceq_i)$ is represented by p and u (and that this representation is unique up to positive affine transformation of u).

III.A Appendix

Lemma III.A.1. *III-A7 holds if and only if, for all C, $\|\Sigma_C\|$ has at most one element (up to positive affine transformation).*

Proof. For any C, if there is at most one element of $\|\Sigma_C\|$, then all simple situations accord the same utilities to consequences (up to positive affine transformation), and thus to constant acts, so III-A7 holds.

On the other hand, if III-A7 holds, then using the von Neumann-Morgenstern theorem on the preferences on constant acts, one obtains a utility function on constant acts, and thus consequences, which applies in all situations. However, in each situation, the utility provided by Theorem III.1 must also represent the preference over constant acts, by (III.3); by the uniqueness properties, such utilities must agree up to positive affine transformation. □

Lemma III.A.2. *III-A9 holds if and only if, for all S, $\|\Xi_S\|$ has at most one element.*

Proof. Take any S, and suppose that there is at most one element of $\|\Xi_S\|$. Consider situations $\sigma_1, \sigma_2 \in Xi_S$ and elements $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ satisfying the conditions of III-A9. The utilities u_1 and u_2 representing the orders in these situations can be scaled so that $u_1(a) = u_2(c)$ and $u_1(b) = u_2(d)$. But then, for any $f \in L(\{a,b\})^S$, $\sum_{S,X_{\sigma_1}} p(s)u_1(x)f(s,x) = \sum_{S,X_{\sigma_2}} p(s)u_2(x)\tau(f)(s,x)$: so the preference orders behave as in III-A9.

Suppose now that III-A9 holds. For any $\sigma_1, \sigma_2 \in \Xi_S$ and any $a, b \in X_{\sigma_1}$ and $c, d \in X_{\sigma_2}$ satisfying the conditions of III-A9, take the representations p_1, u_1 and p_2, u_2 with $u_1(a) = 0, u_1(b) = 1$ and $u_2(c) = 0, u_2(d) = 1$. Consider the following acts in σ_1: for each $A \subseteq S$ f_A, with $f_A(s,b) = 1$ for $s \in A$, $f_A(s,a) = 1$ for $s \notin A$ (and $f(s,x) = 0$ elsewhere); for each element y of $L(\{a,b\})$, the constant act g_y taking the value y. Note that $p_1(A) = \sum_{S,X_{\sigma_1}} p_1(s)u_1(x)f_A(s,x) = \inf_{y \in L(\{a,b\}), f_A \preceq_{\sigma_1} g_y} \sum_{X_{\sigma_1}} y(x)u_1(x)$. However, the image of each of these acts under τ is of the same type (respectively, $f'_A, g'_{y'}$), and u_2 takes the same value as u_1 on the images of the elements involved, so, for all $A \subseteq S$, $p_1(A) = p_2(A)$. □

Proposition III.1. *In the presence of III-A1, III-A10 implies III-A9.*

Proof. Suppose there is more than one element in $\|\Xi_S\|$ for some S. Then there are simple situations $\sigma_1, \sigma_2 \in \Xi_S$, with representations involving p_1 and p_2, where $p_1 \neq p_2$. By III-A10, the situation σ_{12} with set of states S and set of consequences $C_{\sigma_1} \sqcup C_{\sigma_2}$ is simple, and so has representation with probability p_{12}. By III-A1, $\preceq_{\sigma_1}$ coincides with $\preceq_{\sigma_{12}}$ on the acts common to both situations; so by the uniqueness properties of Theorem III.1, they must be represented using the same probability on S. The same goes for $\preceq_{\sigma_2}$: so $p_1 = p_2$, contradicting the supposition. By Lemma III.A.2, III-A9 must hold.

□

Proof of Theorem III.2. Existence. If σ is simple, the existence of the representation (III.4), with $\gamma(s,x) = 1$ for all s, x, is given immediately by Theorem III.1. By III-A7, III-A9 and Lemmas III.A.1 and III.A.2, this representation satisfies the other two clauses.

Suppose σ is not simple. By III-A6, III-A7, III-A8, III-A9 and Lemmas III.A.1 and III.A.2, there is a unique probability function p and a utility function u, unique up to positive affine transformation, satisfying (III.5) and (III.6).

By Postulate III.2, a version of the von Neumann-Morgenstern theorem applies in σ (Fishburn, 1970, p176), giving a representation by an evaluation function $U : S_\sigma \times X_\sigma \to \Re$, unique up to simple positive affine transformation; that is, for all $f, g \in \mathcal{A}_\sigma$,

(III.10) $$f \preceq_\sigma g \text{ iff } \sum_{S_\sigma, X_\sigma} U(s,x).f(s,x) \leqslant \sum_{S_\sigma, X_\sigma} U(s,x).g(s,x)$$

Pick U such that $U(s,x) = 0$ on the null states (if $\preceq_\sigma$ has any).

Define γ the function on $S_\sigma \times X_\sigma$ such that

$$\gamma(s,x) = \begin{cases} \frac{U(s,x)}{p(s).u(x)} & \text{if } p(s) \neq 0 \\ 1 & \text{if } p(s) = 0 \end{cases}$$

By construction and by III-A11, (III.4), (III.5) and (III.6) hold.

Uniqueness. The uniqueness properties of p and u follow from the application of Theorem III.1 to the situations in Ξ_S and Σ_C. Moreover, as noted above, U is given up to simple positive affine transformation. So, for another representation p', u', γ', there are a, b, a' and b'_s for each $s \in S_\sigma$, with a and a' positive, such that, for all $s \in S_\sigma, x \in X_\sigma$, $p'(s) = p(s)$, $u'(x) = a.u(x) + b$ and $p(s).\gamma'(s,x).u'(x) = a'.p(s).\gamma(s,x).u(x) + b'_s$. Substituting for u

and solving for γ', one obtains the required result, with $c = \frac{a'}{a}$ and $d_s = \frac{b_s}{p(s)}$ when $p(s) \neq 0$ and $d_s = 0$ when $p(s) = 0$.

□

Proposition III.2. *III-A13 holds if and only if III-A10 does.*

Proof. Suppose III-A10 but not III-A13. Since X is finite, for each S' partitioning S, there is no infinite ascending chain $X^0 \subset X^1 \subset \ldots$ with $X^j \subseteq X$ for all j and $(S', C^j, \preceq^j)$ simple for all j. So, if III-A13 is not satisfied, there must be a S' partitioning S and X_1 and X_2 with $X_1 \neq X_2$ such that $(S', C_1, \preceq_1)$ and $(S', C_2, \preceq_2)$ are simple, and X_1 and X_2 are maximal. But, by III-A10, $X_{12} = X_1 \cup X_2$ also yields a simple situation $(S', C_{12}, \preceq_{12})$, contradicting the maximality of X_1 and X_2.

Suppose III-A13, and consider any pair of simple small world situations $(S', C_1', \preceq_1')$ and $(S', C_1', \preceq_1')$, for sets of outcomes X_1' and X_2'. Since $X_1', X_2' \subseteq X'$, the maximal set mentioned in III-A13, $X_3' = X_1' \cup X_2' \subseteq X'$, and since $\preceq_3'$ coincides with $\preceq'$ on common acts, $(S', C_3', \preceq_3')$ is simple; so III-A10 holds.

□

Proof of Theorem III.3. The proof proceeds in much the same way as that of Theorem III.2.
Existence. By III-A2-III-A4, a version of the von Neumann-Morgenstern theorem can be applied (Fishburn, 1970, p176), yielding an evaluation function $U : S \times X \rightarrow \Re$ which is unique up to similar positive affine transformation. If $\preceq$ has null states, then pick U such that $U(s, x) = 0$ on the null states. Let u be the restriction of U to constant acts. That is, $u(x) = \sum_{s \in S} U(s, x)$. Note that u is unique up to positive affine transformation.

By III-A12-III-A13 and Theorem III.1, for each S_i, there is a unique $(S_i, X_i, \preceq_i)$ with X_i maximal, and a unique p_i and u_i, unique up to positive affine transformation, representing $\preceq_i$. That is, for all $f, g \in \mathcal{A}_i$,

$$f \preceq_i g \text{ iff } \sum_{s_j \in S_i, x_j \in X_i} p_i(s_j).u_i(x_j).f(s_j, x_j) \leqslant \sum_{s_j \in S_i, x_j \in X_i} p_i(s_j).u_i(x_j).g(s_j, x_j) \quad \text{(III.11)}$$

Since both u_i and u represent the constant acts taking values in C_i, they are positive affine transformations of each other. The orders $\preceq_i$ are thus represented (in the sense of equation (III.11)) by p_i and u.

By III-A12, for each $s \in S$, there exists a unique set $\{s_i^s\}_{i \in I}$, $s_i \in S_i$, with $s = \bigcap_{i \in I} s_i^s$. Define p on S as follows:

(III.12)
$$p(s) = \prod_{i \in I} p_i(s_i^s)$$

Finally, define γ to be the function on $S \times X$ such that

$$\gamma(s,x) = \begin{cases} \frac{U(s,x)}{p(s).u(x)} & \text{if } p(s) \neq 0 \\ 1 & \text{if } p(s) = 0 \end{cases}$$

By construction and by III-A14, p, u and γ satisfy (III.7). Therefore $\sum_{s \in S} p(s).\gamma(s,x).u(x)$ represents the restriction of $\preceq$ to constant acts. By the uniqueness properties of u, as a representation of this order, there is a positive real number α and a real number β such that $\sum_S p(s).\gamma(s,x).u(x) = \alpha u(x) + \beta$ for all $x \in X$. Thus (III.8) is satisfied.

In a similar way, for each $i \in I$, both p, γ and u, and p_i and u represent $\preceq_i$: the first via (III.7), the second with (III.11). But, by construction of p, $p_i(s_i) = \sum_{s \in s_i} p(s) = p(s_i)$. Thus, by the uniqueness the probability and the uniqueness, up to positive affine transformation, of the utility in the representation of $\preceq_i$, there exists a positive real number α_i and a real number β_i such that, for each $s_j \in S_i, x \in X$: $\sum_{s \in s_j} p(s).\gamma(s,x).u(x) = \alpha_i.p_i(s_j).u(x) + \beta_i p_i(s_j) = \alpha_i.p(s_i).u(x) + \beta_i p(s_j)$. So (III.9) is satisfied.

Uniqueness. The uniqueness properties of p follow from Theorem III.1 and III-A13, and the uniqueness properties of u from the von Neumann-Morgestern theorem applied to constant acts. Moreover, as noted above, U is unique up to simple positive affine transformation. So, for another representation p', u', γ', there are a, b, a' and b'_s for each $s \in S$, with a and a' positive, such that, for all $s \in S, x \in X$, $p'(s) = p(s)$, $u'(x) = a.u(x) + b$ and $p(s).\gamma'(s,x).u'(x) = a'.p(s).\gamma(s,x).u(x) + b'_s$. Substituting for u and solving for γ' as in the proof of Theorem III.2, one obtains $\gamma'(s,x) = \frac{a'}{a}\gamma(s,x) - \frac{b.a'}{a}.\frac{\gamma(s,x)}{u'(x)} + \frac{b'_s}{p(s).u'(x)}$, and thus the required result, with $c = \frac{a'}{a}$ and $d_s = \frac{b_s}{p(s)}$ when $p(s) \neq 0$ and $d_s = 0$ when $p(s) = 0$.

□

Bibliography

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Arrow, K. J. (1974). Optimal Insurance and Generalized Deductibles. *Scandinavian Actuarial Journal*, 1:1–42.

Cook, P. J. and Graham, D. A. (1977). The Demand for Insurance and Protection: The Case of Irreplaceable Commodities. *The Quarterly Journal of Economics*, 91:143–156.

Dekel, E., Lipman, B. L., and Rustichini, A. (2001). Representing Preferences with a Unique Subjective State Space. *Econometrica*, 69(4):891–934.

Drèze, J. H. (1987). *Essays on Economic Decisions under Uncertainty*. Cambridge University Press, Cambridge.

Drèze, J. H. and Rustichini, A. (2004). State-Dependent Utility and Decision Theory. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, volume 2. Kluwer, Dordrecht.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Hammond, P. J. (1998). Subjective Expected Utility. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, volume 1. Kluwer, Dordrecht.

Harrison, G. W., List, J. A., and Towe, C. (2007). Naturally Occurring Preferences and Exogenous Laboratory Experiments: A Case Study of Risk Aversion. *Econometrica*, 75:433–458.

Heifetz, A., Meier, M., and Schipper, B. (2006). Interactive Unawareness. *J. Econ. Theory*, 130:78–94.

Hill, B. (2008). Towards a "Sophisticated" Model of Belief Dynamics. Part I: The General Framework. *Studia Logica*, 89(1):81–109.

Hill, B. (2009). Three Analyses of Sour Grapes. In Grüne-Yanoff, T. and Hansson, S. O., editors, *Preference Change: Approaches from Philosophy, Economics and Psychology*. Springer, Theory and Decision Library A, Berlin and New York.

Hill, B. (2010). Awareness Dynamics. *Journal of Philosophical Logic*, 39(2):113–137.

Kahneman, D., Wakker, P. P., and Sarin, R. (1997). Back to {B}entham? {E}xplorations of Experienced Utility. *Quarterly Journal of Economics*, 112:375–405.

Karni, E. (1983). Risk Aversion for State-Dependent Utility Functions: Measurement and Applications. *International Economic Review*, 24:637–647.

Karni, E. (1985). *Decision Making under Uncertainty*. Harvard University Press.

Karni, E. (1993a). A Definition of Subjective Probabilities with State-Dependent Preferences. *Econometrica*, 61:187–198.

Karni, E. (1993b). Subjective expected utility theory with state dependent preferences. *J. Econ. Theory*, 60:428–438.

Karni, E. (1996). Probabilities and Beliefs. *J. Risk Uncertainty*, 13:249–262.

Karni, E. (2007). Foundations of Bayesian Theory. *J. Econ. Theory*, 132:167–188.

Karni, E. (2011). A theory of Bayesian decision making with action-dependent subjective probabilities. *Economic Theory*, 48(1):125–146.

Karni, E. and Mongin, P. (2000). On the Determination of Subjective Probability by Choices. *Management Science*, 46:233–248.

Karni, E. and Schmeidler, D. (1993). On the uniqueness of subjective probabilities. *Economic Theory*, 3:267–277.

Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51:1021–1032.

MacCrimmon, K. R. and Wehrung, D. A. (1990). Characteristics of Risk Taking Executives. *Management Science*, 36:422–435.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Weber, E. U., Blais, A.-R., and Betz, E. (2002). A Domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviours. *Journal of Behavioral Decision Making*, 15:263–290.

IV An additively separable representation in the Savage framework

Abstract

This paper proposes necessary and sufficient conditions for an additively separable representation of preferences in the Savage framework (where the objects of choice are acts: measurable functions from an infinite set of states to a potentially finite set of consequences). A preference relation over acts is represented by the integral over the subset of the product of the state space and the consequence space which corresponds to the act, where this integral is calculated with respect to an evaluation measure on this space. The result requires neither Savage's P3 (monotonicity) nor his P4 (weak comparative probability). Nevertheless, the representation it provides is as useful as Savage's for many economic applications. [1]

Keywords: Expected utility; additive representation; state-dependent utility; monotonicity.

JEL Classification: D81.

---

1. This is a modified version of a paper of same name that appeared in the Journal of Economic Theory, 145(5): 2044-2054, September 2010. The main modifications involve the main result, which contained an error (it claimed countable additivity where only finite additivity can be obtained) and in the introduction of the Constant act-free Monotone Continuity axiom as a condition for countable additivity. Further details available on request.

IV.1 Introduction

According the standard theory of decision under uncertainty, first proposed and axiomatised by Savage (1954), a rational agent's preferences over acts (measurable functions from an infinite set of states to a potentially finite set of consequences) are represented by a probability measure over states and a state-independent utility function on consequences. This representation assumes ordinal state-independence: that if one prefers one consequence to another, one will retain this preference no matter what event is realised. Moreover, it also assumes cardinal state-independence: that one's utility for a particular consequence is the same no matter what event is realised. These assumptions may been challenged.

On the one hand, it has long been recognised that cardinal state-independence does not hold in certain cases: for example, people's risk attitudes many depend on their state of health, and this may affect their choices regarding health and life insurance. On the other hand, although less commonly noted in the literature, ordinal state-independence also seems to fail in many cases, especially when the consequences are commodities. One might prefer a herbal tea at home to a dinner in an expensive restaurant next Saturday if tired, but the dinner to the tea if not. When in ill health, one might prefer living in a built-up area (to be close to the necessary services), whereas in good health, one might prefer the isolation of the countryside. These latter preferences inform one's decisions regarding investment in one's home, which, like health insurance decisions, require one to consider what one's state of health might be in the future and what preferences one would have in different possible states of health.

In such cases, the axioms corresponding to ordinal and cardinal state-independence – monotonicity (Savage's P3) and weak comparative probability (Savage's P4) respectively – are not satisfied. To account for such cases, a weakening of Savage's theory is required; namely, a set of axioms, including *neither* monotonicity *nor* weak comparative probability, which are necessary and sufficient for the following sort of *additively-separable* representation: [2] for all acts f, g,

2. In this formula, we assume that there are finitely many states. This assumption is made entirely for expositional purposes: it is not true in the Savage framework, nor in the result proposed here. See Sections IV.2 and IV.3.

(IV.1) $$f \preccurlyeq g \text{ iff } \sum_{s \in S} U(s, f(s)) \leqslant \sum_{s \in S} U(s, g(s))$$

where U is a real-valued function of states and consequences which we call the *evaluation*. [3] Under the standard theory, and in particular in the presence of the state-independence axioms, this function can be obtained as the product of the probability and the state-independent utility. However, in cases of possible state-dependence, probability and state-independent utility functions do not always exist, and even when they do, they lack economic and empirical meaning because they are not unique in the appropriate sense (see Karni (1996, 2011)). By contrast, the evaluation U exists even in cases of state-dependence and it does have the appropriate uniqueness properties; in cases of state-dependence, it is the more solid economically meaningful concept.

Moreover, it can fill many of the roles traditionally played by probability and utility in economic applications. Consider, for example, the decision maker's choice today about whether to book a table at the restaurant next Saturday or not. To determine what he will choose, only his evaluation U is needed: a factorisation into a probability measure and a utility function adds no useful information. Furthermore, the representation of preferences by an evaluation is entirely sufficient for most standard exercises of comparative statics. For example, the evaluation is all that is needed to determine the effect of changes in the restaurant's prices on the decision maker's choice.

To date, the only axiomatisation of the form (IV.1) in the Savage framework Wakker and Zank (1999) drops weak comparative probability but retains monotonicity, and hence does not apply in the cases mentionned above. In this paper, we provide an axiomatisation in the Savage framework which relies neither on monotonicity nor on weak comparative probability. The theorem is stated in Section IV.2 and is discussed, along with related literature, in Section IV.3. The proof is to be found in the Appendix.

IV.2 Axioms and theorem

Let S be a set of states, with a σ-algebra of events $\mathcal{F}_S$ which contains $\{s\}$ for each $s \in S$, and let C be a set of consequences, with a σ-algebra $\mathcal{F}_C$. C can be finite or infinite:

3. Thanks to Mark Machina and Edi Karni for suggesting this terminology.

if it is finite or countable, $\mathcal{F}_C$ will just be the set of subsets of C; if it is uncountable, this need not be the case. Let $\mathcal{A}$ be the set of measurable functions from S to C – they are called the *acts* – and let $\preccurlyeq$ be a binary relation on $\mathcal{A}$; $\prec$, $\sim$, $\succ$ and $\succcurlyeq$ are defined from $\preccurlyeq$ in the usual way.

Let $\mathcal{A}_p$ be the set of partial measurable functions[4] from S to C; that is, measurable functions from A to C, where $A \subseteq S$. Elements of $\mathcal{A}_p$ are called *partial acts*, and the set A on which a partial act is defined is called its *domain*. Note that, since A can be S, every act belongs to $\mathcal{A}_p$. For any act f and event A, f_A will denote the partial act with domain A which agrees with f on this domain. To each partial act f_A, there corresponds a subset of $S \times C$: namely its graph $\{(s, f_A(s))|\ s \in A\}$. This subset shall also be called f_A. $\mathcal{F}_{SC}$ is defined to be the σ-algebra on $S \times C$ generated by $\mathcal{A}_p$. Note that, when C is finite or countable, $\mathcal{F}_{SC}$ coincides with the product of $\mathcal{F}_S$ and $\mathcal{F}_C$; however, if C is uncountable, this need not be the case.[5]

For A and B disjoint, $f_A g_B$ is the partial function taking the values of f on A and the values of g on B. Given an event A, A^c is the set $S\backslash A$; since it is measurable, it is also an event. As is standard, $\preccurlyeq_A$ will denote the order $\preccurlyeq$ on acts, given the event A. The traditional notion of null event shall be employed: an event A is *null* iff, for any pair of acts $f, g \in \mathcal{A}$, $f \sim_A g$.

We assume two of the basic axioms of Savagean decision theory.

Axiom IV-A1 (Weak order)**.** $\preccurlyeq$ is a weak order: (a) for all f, g in $\mathcal{A}$, $f \preccurlyeq g$ or $g \preccurlyeq f$; and (b) for all f, g and h in $\mathcal{A}$, if $f \preccurlyeq g$ and $g \preccurlyeq h$, then $f \preccurlyeq h$.

Axiom IV-A2 (Sure-thing principle)**.** For any acts f, g, h, h' in $\mathcal{A}$ and any non-null event A, $f_A h_{A^c} \preccurlyeq g_A h_{A^c}$ iff $f_A h'_{A^c} \preccurlyeq g_A h'_{A^c}$

Following Gilboa (1987); Abdellaoui and Wakker (2005), we factorise Savage's P6 axiom into two elements: the standard Archimedean axiom (see for example (Krantz et al., 1971, p204-205)) and solvability (which differs slightly from the solvability axioms pro-

4. Henceforth all functions, partial functions and sets of states shall be assumed to be measurable.

5. Suppose that $S = C = \Re$ with $\mathcal{F}_S$ and $\mathcal{F}_C$ the Borel sets: on the one hand, the graphs of strictly monotonic measurable functions are $\mathcal{F}_{SC}$-measurable but not $\mathcal{F}_S \times \mathcal{F}_C$-measurable; on the other hand, subsets $\{(s, c)|\ s \in (s_1, s_2),\ c \in (c_1, c_2)\}$ for $s_1 < s_2$ and $c_1 < c_2$ are $\mathcal{F}_S \times \mathcal{F}_C$-measurable but not $\mathcal{F}_{SC}$-measurable.

posed in the articles mentioned, in so far as they suppose monotonicity, and this is not assumed here).

Axiom IV-A3 (Archimedean axiom)**.** For each pair of acts f and g, if $A_1, \ldots, A_i, \ldots$ is a sequence of disjoint non-null events such that $f_{A_i} g_{A_i^c} \sim f_{A_j} g_{A_j^c} \succ g$ for all i, j, then it is a finite sequence.

Axiom IV-A4 (Solvability)**.** For acts f, g, h in $\mathcal{A}$ with $f \prec g \prec h$, there exists an event $A \subseteq S$ such that $f_A h_{A^c} \sim g$.

A final technical condition is required, which was first proposed in Hill (2009).

Axiom IV-A5 (Local stability)**.** For any f, g in $\mathcal{A}$, there exists events A_1, A_2 and A_3, such that the non-empty A_i form a partition of S and such that: for any non-null event $B \subseteq A_1$, $f \prec_B g$; for any non-null event $B \subseteq A_2$, $f \sim_B g$; and for any non-null event $B \subseteq A_3$, $f \succ_B g$.

This axiom is required because the set of states is infinite; it is automatically satisfied on finite sets of states.

Theorem IV.1. *For a given S, C, $\preccurlyeq$, suppose that IV-A4 and IV-A5 are satisfied. Then the following are equivalent:*

(i) *$\preccurlyeq$ satisfies IV-A1–IV-A3*

(ii) *there exists a finitely additive measure U on $(S \times C, \mathcal{F}_{SC})$ which takes finite values on $\mathcal{A}$ and is such that, for every $f, g \in \mathcal{A}$,*

$$f \preccurlyeq g \text{ iff } \int_f dU \leqslant \int_g dU \tag{IV.2}$$

Furthermore, let U' be any other finitely additive measure satisfying (IV.2). Then there exists $a > 0$ and a finitely additive measure b on $(S, \mathcal{F}_S)$ such that $U' = aU + b$.

U is the *evaluation* and can be thought of as the analogue of the function U in (IV.1). For further discussion of this representation, see Section IV.3.

Like the probability measure obtained by Savage (1954), U is only finitely additive. It is known, since the work of Villegas (1964) and Arrow (1971), that adding a Monotone

Continuity axiom is necessary and sufficient to ensure countable additivity. The same is true for U, with respect the following version of Monotone Continuity, formulated without using constant acts.

Axiom IV-A6 (Constant act-free Monotone Continuity)**.** For any f, g, h in $\mathcal{A}$ such that $f \succ g$ and sequence of events $\{A_1\}_{i \geqslant 1}$ with $A_1 \supseteq A_2 \supseteq \dots$ and $\cup_{i=1}^{\infty} A_n = \varnothing$, there exists $n \geqslant 1$ such that $h_{A_n} f_{A_n^c} \succ g$ and $f \succ h_{A_n} g_{A_n^c}$.

Proposition IV.1. *For a given S, C, $\preccurlyeq$, suppose that IV-A4 and IV-A5 are satisfied. Then $\preccurlyeq$ satisfies IV-A1–IV-A3 and IV-A6 if and only if it is represented according to (IV.2) by a countably additive measure U on $(S \times C, \mathcal{F}_{SC})$ taking finite values on $\mathcal{A}$.*

The only difference between Constant act-free Monotone Continuity and the standard axiom is the use of a constant act in the place of h. Given that no state-independence axioms are assumed here, constant acts cease to have any particular status, and so it is not surprising that the reference to them in the Monotone Continuity axiom needs to be dropped. Indeed, in the presence of monotonicity (Savage's P3), Constant act-free Monotone Continuity is equivalent to the standard Monotone Continuity axiom.

IV.3 Discussion

Axiomisations for representation (IV.1) have been obtained long ago in frameworks involving a finite number of states and an infinite set of consequences endowed with a rich structure: for example, where the set of consequences is the set of lotteries over a set of outcomes (the so-called Anscombe and Aumann framework) or a connected topological space. However little work has been done on how such a representation can be got in the Savage framework, which assumes a possibly finite set of consequences and an infinite set of states. The main paper in the literature dealing with this problem is Wakker and Zank (1999). We shall bring out the pertinent aspects of Theorem IV.1 via the comparison with their proposal.[6]

The most important difference between the two papers lies in the use of monotonicity: Wakker and Zank need this axiom, whereas Theorem IV.1 does not assume it. As noted in

6. The main points made below also hold for Castagnoli and LiCalzi (2006), which proposes results similar to those in Wakker and Zank (1999), though with a different motivation.

the Introduction, this axiom is violated in several interesting cases, and any full reply to the problem of state-dependent utilities should be able to deal with such cases. The theorem proposed here does apply in such cases, whereas Wakker and Zank's does not.

Another important difference is in the assumptions on the set of consequences. Wakker and Zank require that the set of consequences is a connected topological space, whereas no structure is assumed on the set of consequences here. In particular, Theorem IV.1 applies to finite, as well as infinite, consequence spaces. Not only is this closer to Savage's original theory, where the structural burden is borne by the state space alone, but there are many cases, in particular those involving indivisible goods, where the assumptions on the set of consequences made by Wakker and Zank may be difficult to justify, but where the theorem proposed here continues to apply.

One final difference between the two papers lies in the presentation of the result. It is a non-trivial conceptual problem to find an analogue of (IV.1) when the set of states is infinite (see Wakker and Zank (1999) for an excellent discussion). Wakker and Zank consider two solutions of this problem, both of which are different from, though of course equivalent to, the solution proposed here. To illustrate the relationship, consider Wakker and Zank's second solution (Wakker and Zank, 1999, Theorem 12), which represents preferences by a functional of the form $\int_S u(s, f(s))\, d\mu(s)$, where μ is a non-unique measure on the state space, and u is a function on state-consequence pairs which is not unique up to positive affine transformations. Firstly, the representation proposed here, which involves a single measure U on the Cartesian product of the state space and the consequence space retains a closer analogy with (IV.1), the only difference being the replacement of sums by integrals and a function by a measure. Moreover, given the uniqueness of the evaluation U and the lack of uniqueness of the functions featuring in Wakker and Zank's representation, the former reflects more accurately the economic and empirical content of the representation.

The first two differences mentionned above render Theorem IV.1 more useful in the construction of new theories of state-dependent utilities. The goal of much of the literature on state-dependence is to find necessary and sufficient conditions for the existence of a unique probability function and a suitably unique state-dependent utility function which represents the decision maker's preferences. A popular technique is to assume axioms necessary and sufficient for a representation of the form (IV.1), and then search for new supplementary conditions which guarantee a unique decomposition of the evaluation U

(for example Karni et al. (1983); Karni (2011)). To date, most of this research has had to adopt a framework with a finite state space and a rich consequence space, because, as noted above, these are the only frameworks in which (IV.1) has been axiomatised without relying on monotonicity. Using Theorem IV.1, one can now expect to be able to extend these theories to the Savage framework.

Finally, the result may have useful applications to time preferences. If one interprets S as a set of time points (say, an interval on the real line) and C as a set of consumption bundles, then the objects of choice (acts) are consumption streams of the sort introduced by Strotz (1955). Theorem IV.1 applies, yielding a representation of the form (IV.2). This representation bears the same relationship to the functional form proposed by Strotz as representation (IV.1) bears to the representation proposed by state-dependent utility theorists such as Karni et al. (1983): the single measure on consumption bundle-time point pairs is decomposed into a discount factor and a time-dependent "instantaneous" utility function (Strotz, 1955, p167). And just like for the case of state-dependent utility, all that is required to obtain Strotz's representation are supplementary axioms which guarantee the decomposition. To the knowledge of the author, this is the first axiomatisation in the continuous-time setting of a functional form close to that proposed by Strotz: in particular, contrary to much of the recent literature on time-discounting, it does not suppose that "instantaneous utilities" are time-independent (an assumption which corresponds to state-independence).

IV.A Appendix

Preliminary notions

The proof relies heavily on Theorem 4 in (Krantz et al., 1971, Ch 2). It is worth reproducing the essential definition and the statement of the theorem.

Definition IV.1 (Krantz et al. (1971), p44)**.** Let A be a nonempty set, B and $\gtrsim$ nontrivial binary relations on A and $\circ$ a binary operation from B into A. The quadruple $\langle A, \gtrsim, B, \circ\rangle$ is an *Archimedean, regular, positive, ordered local semigroup* iff, for all $a,\ b,\ c,\ d \in A$, the following eight axioms are satisfied:

1. $\langle A, \gtrsim\rangle$ is a total order: that is, an anti-symmetric weak order (a weak order such that, if $a \geqslant b$ and $b \geqslant a$, then $a = b$).
2. if $(a, b) \in B$, $a \gtrsim c$, and $b \gtrsim d$, then $(c, d) \in B$.
3. if $(c, a) \in B$ and $a \gtrsim b$, then $c \circ a \gtrsim c \circ b$.
4. if $(a, c) \in B$ and $a \gtrsim b$, then $a \circ c \gtrsim b \circ c$.
5. $(a, b) \in B$ and $(a \circ b, c) \in B$ iff $(b, c) \in B$ and $(a, b \circ c) \in B$; and when both conditions hold $(a \circ b) \circ c = a \circ (b \circ c)$.
6. if $(a, b) \in B$, then $a \circ b \succ a$.
7. if $a \succ b$, then there exists $c \in A$ such that $(b, c) \in B$ and $a \gtrsim b \circ c$.
8. $\{n|\ n \in N_a\ and\ b \succ na\}$ is a finite set.

where N_a, a subset of the positive integers, and na, an element of A for each $n \in N_a$, are defined inductively as follows:

(i) $1 \in N_a$ and $1a = a$;

(ii) if $n-1 \in N_a$ and $((n-1)a, a) \in B$, then $n \in N_a$ and na is defined to be $((n-1)a) \circ a$;

(iii) if $n - 1 \in N_a$ and $((n - 1)a, a) \notin B$, then for all $m \geqslant n$, $m \notin N_a$.

The importance of this definition is expressed by the following theorem (Theorem 4 in Krantz et al. (1971)).

Theorem IV.2. *Let $\langle A, \gtrsim, B, \circ\rangle$ be an Archimedean, regular, positive, ordered local semigroup. Then there is a function ϕ from A to $\Re_+$ such that for all $a, b \in A$,* *(i)* $a \succsim b$ *iff* $\phi(a) \geqslant \phi(b)$*;*

(ii) *if* $(a, b) \in B$*, then* $\phi(a \circ b) = \phi(a) + \phi(b)$.

Moreover, if ϕ *and* ϕ' *are two functions from* A *to* $\Re_+$ *satisfying conditions (i) and (ii), then there exists* $\alpha > 0$ *such that, for any nonmaximal* $a \in A$*,* $\phi'(a) = \alpha\phi(a)$.

Proof of Theorem IV.1

The proof employs several propositions and lemmas which are proved separately below. The (ii) to (i) direction is straightforward, so we will restrict attention to the (i) to (ii) direction. Henceforth we assume that IV-A1–IV-A5 hold.

Pick any act $e \in \mathcal{A}$, which shall remain fixed throughout the proof. Let $\mathcal{A}^+ = \{f_A \in \mathcal{A}_p \mid A$ non-null event and for any non-null event $A' \subseteq A, f_{A'}e_{A'^c} \succ e\}$ and $\mathcal{A}^- = \{f_A \in \mathcal{A}_p \mid A$ non-null event and for any non-null event $A' \subseteq A, e \succ f_{A'}e_{A'^c}\}$. Consider the relation $\sim$ on $\mathcal{A}^+$ defined by $f_A \sim g_B$ if $f_A e_{A^c} \sim g_B e_{B^c}$; by A1, this is an equivalence relation. Let $\mathcal{A}^+_\sim$ be the set of equivalence classes of $\mathcal{A}^+$ under $\sim$. $\mathcal{A}^-_\sim$ is defined similarly. $[f_A]$ shall denote the equivalence class (element of $\mathcal{A}^+_\sim$) containing f_A.[7] Define the order $\preccurlyeq^+_\sim$ on $\mathcal{A}^+_\sim$ as follows: $[f_A] \preccurlyeq^+_\sim [g_B]$ iff, for any $f_A \in [f_A]$ and $g_B \in [g_B]$, $f_A e_{A^c} \preccurlyeq g_B e_{B^c}$. Define $\mathcal{B}^+$, a relation on $\mathcal{A}^+_\sim$, as follows: $([f^1_{A_1}], [f^2_{A_2}]) \in \mathcal{B}^+$ iff there exists, for each $i \in \{1, 2\}$, $f^i_{A_i} \in [f^i_{A_i}]$ such that the A_i are pairwise disjoint. Then define $\circ^+$, a binary operation from $\mathcal{B}^+$ into $\mathcal{A}^+_\sim$, as follows: $[f^1_{A_1}] \circ^+ [f^2_{A_2}] = [f^1_{A_1} f^2_{A_2} e_{(A_1 \cup A_2)^c}]$ where for each $i \in \{1, 2\}$, $f^i_{A_i} \in [f^i_{A_i}]$ and the A_i are pairwise disjoint.

Let us first consider the trivial cases. If $\mathcal{A}^+_\sim$ and $\mathcal{A}^-_\sim$ are both empty, then define U to be the measure on $(S \times C, \mathcal{F}_{SC})$ which is 0 everywhere. If one or both of $\mathcal{A}^+_\sim$ and $\mathcal{A}^-_\sim$ are non-empty, but $\mathcal{B}^+$ and $\mathcal{B}^-$ are empty, then $\mathcal{A}^+_\sim$ (resp. $\mathcal{A}^-_\sim$) has one element if it is non-empty (Lemma IV.A.2). In such cases, define the measure U on $(S \times C, \mathcal{F}_{SC})$, by $U(f_A) = 0$ if $f_A e_{A^c} \sim e$, $U(f_A) = 1$ if $f_A e_{A^c} \succ e$ and $U(f_A) = -1$ if $f_A e_{A^c} \prec e$ for any $f_A \in \mathcal{A}_p$. It is straightforward to show in both these cases that U represents $\succcurlyeq$ according to (IV.2).

For the rest of the proof it will be assumed that one or both of $\mathcal{B}^+$ and $\mathcal{B}^-$ are non-empty. The proof for this case rests on the following proposition.

7. So if $f_A \sim g_B$ then $[f_A] = [g_B]$. Often, below, we shall need to chose an arbitrary element from an equivalence class $[f_A]$; for notational convenience we shall often call the element f_A.

Proposition IV.2. *$\langle \mathcal{A}^+_\sim, \succcurlyeq^+_\sim, \mathcal{B}^+, \circ^+ \rangle$ and $\langle \mathcal{A}^-_\sim, \preccurlyeq^-_\sim, \mathcal{B}^-, \circ^- \rangle$ are Archimedean regular positive ordered local semigroups (Definition IV.1).*

It follows from Theorem IV.2 that there is a function $U^+ : \mathcal{A}^+_\sim \to \Re^+$ which respects the order $\succcurlyeq^+_\sim$ (for $[f_A], [g_B] \in \mathcal{A}^+_\sim$, $U^+([f_A]) \geqslant U^+([g_B])$ iff $[f_A] \succcurlyeq^+_\sim [g_B]$) and represents the operation $\circ^+$ by addition (for $[f_A], [g_B] \in \mathcal{B}^+$, $U^+([f_A] \circ^+ [g_B]) = U^+([f_A]) + U^+([g_B])$); moreover, this function is unique up to a positive multiplicative factor. Similarly there is a function $U^- : \mathcal{A}^-_\sim \to \Re^+$ inversing order (for $[f_A], [g_B] \in \mathcal{A}^-_\sim$, $U^-([f_A]) \geqslant U^-([g_B])$) iff $[f_A] \preccurlyeq^-_\sim [g_B]$ and representing $\circ^-$ by addition, which is unique up to a positive multiplicative factor. (Equivalently, $-U^-$ respects the order $\succcurlyeq^-_\sim$ and represents $\circ^-$ by addition.) These naturally induce real-valued functions on $\mathcal{A}^+$ and $\mathcal{A}^-$ sharing the same properties; to ease notation, these functions will also be called U^+ and U^- (so we have $U^+(f_A) = U^+([f_A])$).

It remains to "calibrate" the functions U^+ and U^-; that is, to ensure that the positive and negative utilities add correctly. This only needs to be done if both $\mathcal{A}^+_\sim$ and $\mathcal{A}^-_\sim$ are non-empty; henceforth we suppose that this is indeed the case.

Definition IV.2. $[f_A] \in \mathcal{A}^+_\sim$ and $[g_B] \in \mathcal{A}^-_\sim$ *cancel* if there exist $f_A \in [f_A]$, $g_B \in [g_B]$, such that A and B are disjoint and $f_A g_B e_{(A \cup B)^c} \sim e$.

Let $\mathcal{I}^+ = \{[f_A] \in \mathcal{A}^+_\sim \mid$ there exists $[g_B] \in \mathcal{A}^-_\sim, [f_A]$ and $[g_B]$ cancel$\}$ and let $\mathcal{I}^- = \{[g_B] \in \mathcal{A}^-_\sim \mid$ there exists $[f_A] \in \mathcal{A}^+_\sim, [f_A]$ and $[g_B]$ cancel$\}$. These are non-empty, given that $\mathcal{A}^+_\sim$ and $\mathcal{A}^-_\sim$ are (Lemma IV.A.3). Moreover, there is a natural mapping $\sigma : \mathcal{I}^+ \to \mathcal{I}^-$, taking $[f_A]$ to the $[g_B]$ such that $[f_A]$ and $[g_B]$ cancel. It is straightforward to show that this mapping is well-defined, and that it is one-to-one and onto.

We shall refer to the restrictions of $\mathcal{B}^+$ to $\mathcal{I}^+$ by $\mathcal{B}_\mathcal{I}{}^+$ and the restrictions of $\succcurlyeq^+_\sim$ and $\circ^+$ by $\succcurlyeq^+_\sim$ and $\circ^+$ respectively; similarly for $\mathcal{I}^-$. Note that $\mathcal{B}_\mathcal{I}{}^+$ is non-empty if and only if $\mathcal{B}_\mathcal{I}{}^-$ is, so there are two cases.

Case 1. $\mathcal{B}_\mathcal{I}{}^+$ and $\mathcal{B}_\mathcal{I}{}^-$ are empty. In this case, $\mathcal{I}^+$ and $\mathcal{I}^-$ contain one element each (the proof of this is a simple extension of the proof of Lemma IV.A.2); let the elements be $\widehat{[f_A]}$ and $\sigma(\widehat{[f_A]})$ respectively. So for $[g_B] \in \mathcal{I}^-$, $U^-([g_B]) = \alpha U^+(\sigma^{-1}([g_B]))$, where $\alpha = \frac{U^-(\sigma(\widehat{[f_A]}))}{U^+(\widehat{[f_A]})}$.

Case 2. $\mathcal{B}_\mathcal{I}{}^+$ and $\mathcal{B}_\mathcal{I}{}^-$ are non-empty. By inspection of the proof of Proposition IV.2, it is easily seen that $\langle \mathcal{I}^+, \succcurlyeq^+_\sim, \mathcal{B}_\mathcal{I}{}^+, \circ^+ \rangle$ and $\langle \mathcal{I}^-, \preccurlyeq^-_\sim, \mathcal{B}_\mathcal{I}{}^-, \circ^- \rangle$ are Archimedean regular positive ordered local subgroups, so Theorem IV.2 applies. U^- represents the restriction of $\preccurlyeq^-_\sim$ to $\mathcal{I}^-$; however, $U^+.\sigma^{-1}$ does as well.

Proposition IV.3. *For* $[g_B], [g'_{B'}] \in \mathcal{I}^-$, $U^+(\sigma^{-1}([g_B])) \geqslant U^+(\sigma^{-1}([g'_{B'}]))$ *iff* $[g_B] \preccurlyeq^-_\sim [g'_{B'}]$.

Since, by Theorem IV.2, any two representations of the restriction of $\preccurlyeq^-_\sim$ to $\mathcal{I}^-$ are related by a positive multiplicative transformation, there exists an $\alpha > 0$ such that, for all $[g_B] \in \mathcal{I}^-$, $U^-([g_B]) = \alpha U^+(\sigma^{-1}([g_B]))$.

Define $U : \mathcal{A}_p \to \Re$ as follows. Firstly define it on $\mathcal{A}^+ \cup \mathcal{A}^- \cup \{f_A \in \mathcal{A}_p \mid \text{ for any event } A' \subseteq A, f_{A'}e_{A'^c} \sim e\}$:

$$U([f_A]) = \begin{cases} U^+(f_A) & \text{if } f_{A'}e_{A'^c} \succ e \text{ for all non-null events } A' \subseteq A \\ -\frac{1}{\alpha}U^-([f_A]) & \text{if } f_{A'}e_{A'^c} \prec e \text{ for all non-null events } A' \subseteq A \\ 0 & \text{if } f_{A'}e_{A'^c} \sim e \text{ for all events } A' \subseteq A \end{cases} \tag{IV.3}$$

For any $f_A \in \mathcal{A}_p$, IV-A5 applied to $f_A e_{A^c}$ and e ensures that there is a partition of A into three events A_1, A_2, A_3 such that, on each event, one of the three conditions in (IV.3) holds. So U can be defined for any $f_A \in \mathcal{A}_p$ by $U(f_A) = U(f_{A_1}) + U(f_{A_2}) + U(f_{A_3})$.

Given that $\mathcal{F}_{SC}$ is the σ-algebra generated by $\mathcal{A}_p$, this suffices to define U on $\mathcal{F}_{SC}$. By construction, U is a finitely additive measure on $(S \times C, \mathcal{F}_{SC})$, which is finite-valued on $\mathcal{A}$ and represents $\succcurlyeq$ according to (IV.2).

Uniqueness. Assume that $\succcurlyeq$ is non-trivial and let U' be any other measure on $(S \times C, \mathcal{F}_{SC})$ taking finite values on $\mathcal{A}$ and representing $\succcurlyeq$. Define $b : \mathcal{F}_S \to \Re$ by $b(A) = U'(e_A)$, where e is the act used in the construction of U (e is such that, for any event A, $U(e_A) = 0$). By the additivity properties of U', b is a finite-valued finitely additive measure on $(S, \mathcal{F}_S)$. Consider the measure $U' - b.p_1$ on $(S \times C, \mathcal{F}_{SC})$, where p_1 is the projection onto S: $p_1(f_A) = A$. For any $f_A \in \mathcal{A}^+$ and for all non-null $A' \subseteq A$, $f_{A'} \succ e_{A'}$, so $U'(f_{A'}) > U'(e_{A'})$ and hence $(U' - b.p_1)(f_{A'}) > 0$. $U' - b.p_1$ is thus strictly positive on $\mathcal{A}^+$. Furthermore, $U' - b.p_1$ is well-defined on the set of equivalence classes $\mathcal{A}^+_\sim$; we now show that $U' - b.p_1$ represents $\succcurlyeq^+_\sim$. Let $[f_A], [g_B] \in \mathcal{A}^+_\sim$ and suppose that $[f_A] \succcurlyeq^+_\sim [g_B]$. By Lemma IV.A.1, there exists non-null $A' \subseteq A$ such that $f_{A'} \in [g_B]$. But $(U' - b.p_1)([f_A]) - (U' - b.p_1)([g_B]) = (U' - b.p_1)(f_A) + (U' - b.p_1)(e_{A^c}) - (U' - b.p_1)(f_{A'}) - (U' - b \circ p_1)(e_{A'^c}) = U'(f_{A \backslash A'}) - U'(e_{A \backslash A'}) \geqslant 0$ since $f_{A \backslash A'} e_{(A \backslash A')^c} \succcurlyeq e$. Similarly, if $(U' - b.p_1)([f_A]) - (U' - b.p_1)([g_B]) \geqslant 0$ then $U'(f_A e_{A^c}) \geqslant U'(g_B e_{B^c})$ and hence $[f_A] \succcurlyeq^+_\sim [g_B]$, since U' represents $\succcurlyeq$. So $U' - b.p_1$ represents $\succcurlyeq^+_\sim$; a similar argument holds for $\preccurlyeq^-_\sim$ and $\mathcal{A}^-_\sim$. By the uniqueness clause in Theorem IV.2, $(U' - b \circ p_1) = a.U$ for some $a > 0$.

This completes the proof of Theorem IV.1.

Proof of Proposition IV.1

The right-to-left entailment is straightforward. As concerns the left-to-right direction, due to a well-known result, it suffices to show that for any $f \in \mathcal{A}$ and sequence $A_i \in \mathcal{F}_S$ with $A_i \downarrow \varnothing$, $U(f_{A_i}) \to 0$ for U representing $\preccurlyeq$ according to (IV.2). By IV-A5, to establish this, it suffices to show that, for any $A_i \downarrow \varnothing$ and $f_{A_1} \in \mathcal{A}^+$, either there exists $n \geqslant 1$ such that A_i are null for all $i \geqslant i$, or $U^+(f_{A_i}) \to 0$ (where $\mathcal{A}^+$ and U^+ are as defined in the proof of Theorem IV.1 above); and likewise for U^-. We consider the case of U^+; U^- is treated similarly. Consider any such $f_{A_1} \in \mathcal{A}^+$ and $A_i \downarrow \varnothing$. If there exists $n \geqslant 1$ with A_i null for all $i \geqslant n$, then the result is immediate, so suppose that this is not the case.

As is clear from the proof of Theorem IV.2 (Krantz et al., 1971, §2.2.2), we can distinguish two cases concerning its application to the derivation of U^+ and U^- in the proof of Theorem IV.1 (see in particular Proposition IV.2). Either there is a smallest element $[h_C]$ of $\mathcal{A}^+_\sim$ (for all $[f_A] \in \mathcal{A}^+_\sim$, $[h_C] \preccurlyeq^+_\sim [f_A]$), and every element of $\mathcal{A}^+_\sim$ is of the form $n[h_C]$ for some positive integer n (Krantz et al., 1971, §2.2.2), or there is no smallest element of $\mathcal{A}^+_\sim$. In the first case there are no infinite decreasing sequences of elements of $\mathcal{A}^+_\sim$, whereas in the latter there are.

If $\mathcal{A}^+_\sim$ has a smallest element, then there exists $n \geqslant 1$ such that A_i is null for all $i \geqslant n$, contradicting the assumption above. We now consider the case where $\mathcal{A}^+_\sim$ has no smallest element. By Lemma IV.A.1 and the fact that there exist infinite decreasing sequences of elements of $\mathcal{A}^+$, there exists $n \geqslant 1$ and $g_B \in \mathcal{A}^+$ such that B is disjoint from A_n, and a sequence of non-null B_k, $k \geqslant 1$ such that $g_{B_k} \succcurlyeq g_{B_{k+1}}$, $B_k \subseteq B$ for all $k \geqslant 1$, and $U^+(g_{B_k}) \to 0$ as $k \to \infty$. For each such k, since $g_{B_k} e_{B_k^c} \succ e$ (where e is as stated in the proof of Theorem IV.1), by IV-A6 there exists $i \geqslant N$ such that $g_{B_k} e_{B_k^c} \succ f_{A_i} e_{A_i^c}$. If follows from the fact that U^+ is order preserving (Proposition IV.2 and Theorem IV.2) that $U^+(g_{B_k}) > U^+(f_{A_i}) \geqslant \lim_{i\to\infty} U^+(f_{A_i})$. Taking the limit as $k \to \infty$ it follows that $\lim_{i\to\infty} U^+(f_{A_i}) = 0$. So U is countably additive, as required.

Proofs of Auxiliary Propositions and Lemmas

The following lemmas will be used below or have been mentioned above.

Lemma IV.A.1. *For $f_A, g_B \in \mathcal{A}^+$, if $[f_A] \succsim^+_\sim [g_B]$, then there exists a non-null event $A' \subseteq A$, such that $f_{A'} \in [g_B]$.*

Proof. If $f_A e_{A^c} \sim g_B e_{B^c}$, let $A' = A$. If not, $f_A e_{A^c} \succ g_B e_{B^c} \succ e$, so, by IV-A4, there is an event $A'' \subseteq A$ such that $f_{A \backslash A''} e_{A^c \cup A''} \sim g_B e_{B^c}$. Setting $A' = A \backslash A''$ yields the required result. If A' were null, then $g_B e_{B^c} \sim f_{A'} e_{A'^c} \sim e$, contradicting $g_B \in \mathcal{A}_+$. □

Lemma IV.A.2. *If $\mathcal{A}^+_\sim$ (resp. $\mathcal{A}^-_\sim$) is non-empty and $\mathcal{B}^+$ (resp. $\mathcal{B}^+$) is empty, then $\mathcal{A}^+_\sim$ (resp. $\mathcal{A}^-_\sim$) has just one element.*

Proof. Suppose not, and take $f_A, g_B \in \mathcal{A}^+$ with $[f_A] \succ^+_\sim [g_B]$. By Lemma IV.A.1, there is a non-null event $A' \subseteq A$ such that $f_{A'} \in [g_B]$, but by IV-A2 and the fact that $[f_A] \neq [g_B]$, $A \backslash A'$ must also be non null. $([f_{A'}], [f_{A \backslash A'}]) \in \mathcal{B}^+$, contradicting the assumption that it was empty. □

Lemma IV.A.3. *If $\mathcal{A}^+_\sim$ and $\mathcal{A}^-_\sim$ are non-empty, then $\mathcal{I}^+$ and $\mathcal{I}^-$ are non-empty.*

Proof. Consider arbitrary $f_A \in \mathcal{A}^+$ and $g_B \in \mathcal{A}^-$. If A and B have non-null intersection, then Axiom IV-A4 applied to $f_A e_{A^c} \succ e \succ g_B e_{B^c}$ yields $f_{A'} g_{B'} e_{(A' \cup B')^c} \sim e$, where $A' \subseteq A$, $B' \subseteq B$ are non-null events; so $[f_{A'}]$ and $[g_{B'}]$ cancel. If A and B have null intersection then they can be assumed to be disjoint (remove the intersection from one). Consider $f_A g_B e_{(A \cup B)^c}$: if $f_A g_B e_{(A \cup B)^c} \sim e$ then $[f_A]$ and $[g_B]$ cancel; if $f_A g_B e_{(A \cup B)^c} \succ e$, then applying IV-A4 to $f_A g_B e_{(A \cup B)^c} \succ e \succ g_B e_{B^c}$ gives $f_{A'} g_{B'} e_{(A' \cup B')^c} \sim e$ where $A' \subseteq A$ and $B' \subseteq B$ so $[f_{A'}]$ and $[g_{B'}]$ cancel; similarly for $f_A g_B e_{(A \cup B)^c} \prec e$. □

Proof of Proposition IV.2. Since the cases are similar, we shall only treat the case of $\langle \mathcal{A}^+_\sim, \succsim^+_\sim, \mathcal{B}^+, \circ^+ \rangle$ here. We shall show that the clauses of Definition IV.1 are satisfied.

1. $\precsim^+_\sim$ *is a total order.* The order $\precsim^+_\sim$ on $\mathcal{A}^+_\sim$ inherits the properties of connectedness and transitivity from the order $\precsim$ on $\mathcal{A}$: so axiom IV-A1 guarantees that $\precsim^+_\sim$ is a weak order. Furthermore, since $\mathcal{A}^+_\sim$ is obtained by quotienting by $\sim$, $\precsim^+_\sim$ is anti-symmetric. It is thus a total order.

2. *If $([f_A], [g_B]) \in \mathcal{B}^+$, $[f_A] \succsim^+_\sim [f'_{A'}]$ and $[g_B] \succsim^+_\sim [g'_{B'}]$, then $([f'_{A'}], [g'_{B'}]) \in \mathcal{B}^+$.* Since $([f_A], [g_B]) \in \mathcal{B}^+$, there are elements $f_A \in [f_A]$, $g_B \in [g_B]$, with A and B disjoint. Using

Lemma IV.A.1, choose events $A'' \subseteq A$ and $B'' \subseteq B$ such that $f_{A''} \in [f'_{A'}]$ and $g_{B''} \in [g'_{B'}]$. Since A and B are disjoint, so are A'' and B'', and hence $([f'_{A'}], [g'_{B'}]) \in \mathcal{B}^+$.

3. and 4. Note that, since $\circ^+$ is commutative, Clause 3. is satisfied if and only if Clause 4. is. So it suffices to establish that: *If* $([f_A], [g_B]) \in \mathcal{B}^+$ *and* $[f_A] \succsim^+ [h_C]$, *then* $[f_A] \circ^+ [g_B] \succsim^+ [h_C] \circ^+ [g_B]$. Since $([f_A], [g_B]) \in \mathcal{B}^+$, there are elements $f_A \in [f_A]$, $g_B \in [g_B]$, with A and B disjoint. Using Lemma IV.A.1, choose an event $A' \subseteq A$ such that $f_{A'} \in [h_C]$. By the definition of $\mathcal{A}^+_{\sim}$, $f_{A'}e_{A'^c} \preccurlyeq f_A e_{A^c}$; by Axiom IV-A2, it follows that $f_{A'}g_B e_{(A' \cup B)^c} \preccurlyeq f_A g_B e_{(A \cup B)^c}$. Hence $[h_C] \circ^+ [g_B] \precsim^+ [f_A] \circ^+ [g_B]$.

5. $([f_A], [g_B]) \in \mathcal{B}^+$ *and* $([f_A] \circ^+ [g_B], [h_C]) \in \mathcal{B}^+$ *iff* $([g_B], [h_C]) \in \mathcal{B}^+$ *and* $([f_A], [g_B] \circ^+ [h_C]) \in \mathcal{B}^+$*; and when both conditions hold* $([f_A] \circ^+ [g_B]) \circ^+ [h_C] = [f_A] \circ^+ ([g_B] \circ^+ [h_C])$. If $([f_A], [g_B]) \in \mathcal{B}^+$ and $([f_A] \circ^+ [g_B], [h_C]) \in \mathcal{B}^+$, then there are f_A, g_B and h_C, members of $[f_A]$, $[g_B]$ and $[h_C]$ respectively, such that A, B and C are disjoint. It follows that $([g_B], [h_C]) \in \mathcal{B}^+$ and $([f_A], [g_B] \circ^+ [h_C]) \in \mathcal{B}^+$; furthermore $([f_A] \circ^+ [g_B]) \circ^+ [h_C] = [f_A g_B h_C] = [f_A] \circ^+ ([g_B] \circ^+ [h_C])$. The same argument works in the other direction.

6. *If* $([f_A], [g_B]) \in \mathcal{B}^+$, *then* $[f_A] \circ^+ [g_B] \succ^+_{\sim} [f_A]$. Since $([f_A], [g_B]) \in \mathcal{B}^+$, there are elements $f_A \in [f_A]$, $g_B \in [g_B]$, with A and B disjoint. $f_A g_B e_{(A \cup B)^c}$ and $f_A e_{A^c}$ differ solely on B; by IV-A2, it is their comparison on this set that decides the preference ordering between them. Furthermore, by definition of $\mathcal{A}^+$, $g_B e_{B^c} \succ e$; hence the required result.

7. *If* $[f_A] \succ^+_{\sim} [g_B]$, *then there exists a* $[h_C] \in \mathcal{A}^+_{\sim}$ *with* $([g_B], [h_C]) \in \mathcal{B}^+$ *and* $[f_A] \succsim^+ [g_B] \circ^+ [h_C]$. By Lemma IV.A.1, there is an $A' \subseteq A$ such that $f_{A'} \in [g_B]$. Since $f_A e_{A^c} \succ f_{A'} e_{A'^c}$, $A \backslash A'$ is a non-null event; take C to be an arbitrary non-null subevent of $A \backslash A'$. Since $f_A \in \mathcal{A}^+$ and C is non-null, $f_C \in \mathcal{A}^+$; since $A' \cup C \subseteq A$, $f_{A' \cup C} e_{(A' \cup C)^c} \preccurlyeq f_A e_{A^c}$. Taking $[h_C] = [f_C]$ gives the result.

8. *For all* $[f_A], [g_B] \in \mathcal{A}^+_{\sim}$, $\{n \mid n \in N$ *and* $[g_B] \succ n[f_A]\}$ *is finite, where* $N_{[f_A]}$ *and* $n[f_A]$ *as in Definition IV.1*. Suppose not. Then there exists an infinite sequence $f^i_{A_i} \in \mathcal{A}^+$ such that for all i, j, $f^i_{A_i} e_{A_i^c} \sim f^j_{A_j} e_{A_j^c} \succ e$ and the A_i are disjoint. The acts $f = f^1_{A_1} \dots f^i_{A_i} \dots e_{(\bigcup A_i)^c}$ and e, and the infinite sequence of disjoint non-null events A_i violate IV-A3.

This concludes the proof of Proposition IV.2. □

Proof of Proposition IV.3. Suppose not: there are $[g_B], [g'_{B'}] \in \mathcal{I}^-$ with $U^+(\sigma^{-1}([g'_{B'}])) \geqslant U^+(\sigma^{-1}([g_B]))$ but $[g'_{B'}] \succ^-_{\sim} [g_B]$. Pick $f_A \in \sigma^{-1}([g'_{B'}])$ and let $A'' \subseteq A$ be such that

$f_{A''} \in \sigma^{-1}([g_B])$ (such an A'' exists by Lemma IV.A.1). Similarly, pick $g_B \in [g_B]$ with B disjoint from A, and let $B'' \subset B$ be such that $g_{B''} \in [g'_{B'}]$. Since $[g'_{B'}] \neq [g_B]$, $B \backslash B''$ is non null. By definition, $f_A g_{B''} e_{(A \cup B')^c} \sim e \sim f_{A''} g_B e_{(A'' \cup B)^c}$, and $f_A \succcurlyeq f_{A''}$. By IV-A2, it follows that $g_{B \backslash B''} e_{(B \backslash B'')^c} \sim f_{A \backslash A''} e_{(A \backslash A'')^c} \succcurlyeq e$, contradicting the fact that $g_B \in \mathcal{A}^-$. The assumption that $U^+(\sigma^{-1}([g'_{B'}])) \geqslant U^+(\sigma^{-1}([g_B]))$ but $[g'_{B'}] \succsim^- [g_B]$ is thus false. □

Bibliography

Abdellaoui, M. and Wakker, P. (2005). The Likelihood Method for Decision under Uncertainty. *Theory Dec.*, 58:3–76.

Arrow, K. J. (1971). *Essays in the Theory of Risk Bearing*. Markham Publishing., Chicago.

Castagnoli, E. and LiCalzi, M. (2006). Benchmarking real-valued acts. *Games Econ. Behav.*, 57:236–253.

Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *J. Math. Econ.*, 16:65–88.

Hill, B. (2009). When is there state independence? *J. Econ. Theory*, 144(3):1119–1134.

Karni, E. (1996). Probabilities and Beliefs. *J. Risk Uncertainty*, 13:249–262.

Karni, E. (2011). A theory of Bayesian decision making with action-dependent subjective probabilities. *Economic Theory*, 48(1):125–146.

Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51:1021–1032.

Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, volume 1. Academic Press, San Diego.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Strotz, R. H. (1955). Myopia and Inconsistency in Dynamic Utility Maximization. *Rev. Econ. Stud.*, 23(3):165–180.

Villegas, C. (1964). On Qualitative Probability/sigma-Algebras. *The Annals of Mathematical Statistics*, pages 1787–1796.

Wakker, P. and Zank, H. (1999). State Dependent Expected Utility for Savage's State Space. *Mathematics of Operations Research*, 24:8–34.

V Confidence and decision

Abstract

Many real-life decisions have to be taken on the basis of probability judgements of which the decision maker is not entirely sure. This paper develops a decision rule for taking such decisions, which incorporates the decision maker's confidence in his probability judgements according to the following maxim: the larger the stakes involved in a decision, the more confidence is required in a probability judgement for it to play a role in the decision. A formal representation of the decision maker's confidence is proposed and used to formulate a family of decision models conforming to this maxim. A natural member of this family is studied in detail. It is structurally simpler than other recent models of decision under uncertainty, which may make it easier to apply to practical decisions, whilst being axiomatically sound, permitting the separation of beliefs and tastes, and allowing comparative statics analysis of attitudes to choosing in the absence of confidence. [1]

Keywords: Confidence; multiple priors; confidence ranking; cautiousness coefficient; ambiguity.

JEL classification: D81, D83.

1. This is a slightly modified version of a paper of same name that appeared in Games and Economic Behavior, 82: 675-692, November 2013

V.1 Introduction

V.1.1 General motivation

A regional governor is faced with the decision whether to permit the construction of a factory in his area. To evaluate a major uncertainty in the choice – the possibility of damage to the district farming area (which specialises in maize) – he commissions meteorologists to estimate the probability that the gases emitted by the proposed factory, which are known to damage crops, reach the farming area. Although there is a fair amount of disagreement among the experts, the governor considers that a value of at most 10^{-5} for the probability is quite representative of the opinions. On the basis of this probability estimate and his utilities for the relevant consequences, which he determines using standard decision-analysis techniques, he decides to grant permission for the construction. However, the project falls through, and a different project comes before him for the same land, namely to use it for GM maize crops (which are wind pollinated). Evaluating the possible consequences as before, it becomes clear that, while many aspects are analogous to those for the factory, the orders of magnitude are far greater: both the advantages of allowing the development and its main potential undesirable consequence – cross-pollination of GM crops with non-GM crops – are larger by a factor of a thousand, in the governor's opinion. As concerns the probability that pollen from the GM crops arrives at the non-GM farming area, the experts guarantee that this is the same as the probability of fumes from the factory arriving at the maize fields, since both depend on the same meteorological factors. Were the governor to decide as before, using his utilities and the 10^{-5} estimate adopted previously, he would approve the GM project. However, he is uncomfortable with this choice. Although he was happy to rely on the probability estimate of at most 10^{-5} for the decision involving the factory, he is not sure enough in this estimate when it comes to deciding about the GM crops: after all, given the potentially grave consequences of GM infection, the stakes are higher in the latter decision.

At first glance, it may seem that the governor is violating one of the tenets of decision theory, namely the separation of beliefs and tastes. The standard normative theory of decision prohibits altering the beliefs, or probability judgements, used depending on the decision faced. However, this theory ignores the fact that decision makers may be more or less *confident* in the beliefs which inform their decisions, and that confidence in beliefs is potentially relevant for decision making. Rather than using different beliefs in the two decisions, the governor could be understood as demanding different levels of confidence in the beliefs he uses. When the stakes are mediumly high – at worst the agricultural yield suffers for a (relatively) short period due to pollution – he allows himself to rely on a judgement – that the probability is at most 10^{-5} – which he is not confident enough in to invoke when the stakes are higher – when the worst that can happen is that the GM genes infect the non-GM population. In the light of these considerations, the governor ceases to appear irrational. Indeed, he can be understood as invoking the following, perfectly reasonable maxim: the more important the decision, the more confidence is required in a belief for it to play a role in that decision.

The main aim of this paper is to develop a decision rule based on this maxim. This rule should be normatively plausible, as well as simple and tractable, and hence applicable to decisions such as the one considered above. Such a decision rule would be able to incorporate the intuitions just mooted, without violating general principles of rationality, such as the separation of beliefs and tastes.

A decision rule of this sort may be useful for prescriptive purposes in several domains. In decisions about policies to adopt in the face of climate change, for example, the stakes are high and the probabilities are uncertain; the same goes for other domains where the Precautionary Principle has been an object of debate.[2] A decision rule of the sort considered here naturally applies in such cases, and recommends only relying on probability judgements in which one has enough confidence to match the gravity of the potential consequences.

Moreover, the rule may have implications for the use made of conclusions of statistical studies (in classical statistics). Often, confidence intervals at a standard fixed level (for example, the 95% level) are used to inform decisions. The proposed decision rule would suggest that one may vary the confidence intervals one uses in decisions: whereas if relatively little is at stake one may use confidence intervals at, say, the 90% level, if there is a lot at stake one might insist on relying only on intervals that are at, say, the 99% level.

Finally, the theory developed below may be descriptively relevant. Consider the case of an investor who invests in a start-up drug firm whose major product is a new vaccine

2. One specificity of such decisions, which plays a role in many analyses, is the irreversability of the consequences of actions. This of course implies that the stakes are higher than they otherwise would be: one cannot undo tomorrow what has been done today.

which has just come out on the market amid some controversy, though he refuses to let his daughter take the vaccine. This is naturally taken account of under the proposed rule: the investor is confident enough in his judgement about the probability that the drug will be a success to take the investment decision on the basis of that judgement, but not confident enough to rely on it in the more important personal decision.

In this paper, we propose a formal model of a decision maker's confidence in his beliefs and use it to develop a decision rule that adheres to the maxim that the higher the stakes involved in a decision, the more confidence is required in a probability judgement for it to play a role in that decision. In reality, a family of such decision rules can be defined on the basis of the model of confidence proposed, and the main purpose of the paper is to promote this family. To this end, and for the sake of concreteness, we perform an in-depth axiomatic study of a single, natural member of the family. Perhaps surprisingly given the simplicity of the model and the maxim on which it is based, it corresponds to relatively attractive properties of preferences, which are essentially stakes-corrected versions of standard axioms. Moreover, it provides a full separation of utilities, belief attitudes (including confidence in beliefs) and attitude to choosing in the absence of confidence. Relations to the standard criteria for comparative ambiguity aversion are also examined. Prior to the formal analysis, we further motivate the need for a new model to capture the role of confidence in beliefs in decision by relating it to a well-known behavioural axiom, and present informally the main aspects of our proposal, as well as connections with existing literature.

V.1.2 Confidence and behaviour: an example

The relevance of confidence in beliefs for decision making in accordance with the maxim mentioned above is particularly evident behaviourally in certain violations of Gilboa and Schmeidler's (1989) C-independence axiom. Under the assumption of linear utility for money, this axiom basically states that preferences between a pair of acts (that is, functions from states to consequences) are unaffected by multiplying the values of the consequences delivered by the acts by a positive number and by adding a constant to them. In order to see the relevance of this property for confidence, let us consider a precise example that can be thought as a stylized version of the governor's decision described at the beginning of the paper. [3]

Suppose that the governor has linear utility, [4] and consider an urn in which he is told that there are one million balls, each of which is coloured either red or blue. He knows for sure that at least 990 000 balls in the urn are blue, and that at least one ball in the urn is red. Moreover, his advisers (who are experts on urns, but have not been able to count the balls in this urn) estimate that at most ten of the balls in the urn are red. Given this situation, consider how he prices the bet f given in Figure V.1 – that is, consider the value for which he is indifferent between receiving that value for sure and playing the bet. Suppose for example that he prices the bet at $-\$0.1M$ (this is given in Figure V.1 in the line p_f). It follows from C-independence that the governor should be indifferent between $\alpha f + \eta$ and $\alpha p_f + \eta$ in Figure V.1 for any positive α and any η. [5] For example, to satisfy this axiom, he must price the bet g, which yields $\$10000$ if the ball is blue and $-\$1M$ if the ball is red, at $-\$100$ (take $\alpha = 10^{-3}$ and $\eta = 0$). More generally, C-independence implies that whenever the governor gives f a negative price (and so prefers p_0 in Figure V.1, which yields $\$0$ for sure, to f) then he must do the same for the bet g. This is at odds with the behavioural pattern described in the previous section. The urn could be thought of as a stylized representation of the meteorological conditions and the governor's information concerning them, the bet g could be considered as a stylized version of the option of allowing the construction of the factory, and the bet f could be considered as analogous to the option of allowing planting of GM crops. As noted previously, it does not seem unreasonable for the governor to reject the GM project (ie. to prefer p_0 to f), whilst accepting the construction of the factory (ie. preferring g to p_0), and this is incompatible with C-independence.

Since the governor's pattern of choice in this example violates C-independence, it cannot be captured by decision models satisfying this axiom, such as the standard expected

3. Formal definitions and technical details in a general framework, as well as related discussion, are given in Section V.2.

4. Of course, this assumption is not essential to the example, but is made for ease of presentation. To dispense with it, it suffices to interpret the values of consequences presented below as given in utiles, or consider that the ball drawn from the urn determines the composition of a second urn in such a way that the utility values of the second lottery coincide with those given.

5. Under the assumption of linear utility, these acts correspond to mixture of f (resp. p) with $\frac{\eta}{1-\alpha}$, in the sense that will be defined in Section V.2.1.

utility model and the maxmin expected utility model of Gilboa and Schmeidler (1989).[6] Indeed, for the purposes of gaining intuition, it may be useful to consider the choices in the light of the latter model, which evaluates an act by the lowest expected utility calculated over a fixed set of probability measures. In the example considered, the lowest expected utility is attained with the highest probability of the 'bad' event, namely the drawing of a red ball. The indifference between f and p_f is consistent (under the assumption that utility is linear) with the use of 0.01 as the highest possible probability of getting a red ball. So, in pricing the bet this way, the governor takes the worst probability out of those that, for all he knows, may be relevant; in doing so, he ignores his advisers' estimates. (Recall that he knows for sure that there are at most 10 000 red balls in the urn.) If, on the other hand, he gives g a positive price, this would correspond in the maxmin expected utility model to him taking as worst-case probability of drawing a red ball a value which is lower than 0.01 (and more precisely, lower than $\frac{10}{1010}$). This does not seem unreasonable: indeed, if he relied on his advisers' estimates in that decision, the worst-case probability would be 10^{-5}, and he would price the bet at $\$9989.90$. Hence, although the maxmin expected utility representation does not allow the use of different worst-case probabilities for different decisions, one can understand why one might want to use them: in the high-stakes decision (concerning f), one is not confident enough in the advisers' estimates of the worst-case probability to rely on them, whereas in the lower-stakes decision (concerning g), one is confident enough in the estimates to use them, and hence to work with a lower worst-case probability for the 'bad' event.

Although some more recent decision models do not satisfy C-independence, it is not clear that they can capture the sorts of violations that appear to be related to confidence as just described. To illustrate this, consider the popular smooth ambiguity model proposed by Klibanoff et al. (2005) and others (see Section V.1.4), and the following extension of the example given above. Suppose that, beyond the preferences detailed above, the governor is indifferent between $\alpha f + \eta$ and $\alpha p_f + \eta$ whenever the worst consequence of the bet $\alpha f + \eta$ is below $-\$10M$. As noted, this can be understood as him ignoring his advisers' estimates and evaluating the bet using the worst probability value out of those that may be relevant given what he knows for sure about the urn. Moreover, such preferences

6. A large class of models satisfying this axiom – those of invariant biseparable preferences – is studied in Ghirardato et al. (2004). As they note, C-independence essentially corresponds to a property of functional representations called constant linearity.

Figure V.1 – Bets (values in dollars, ‘M’ stands for ‘million’)

	Colour of ball drawn from urn	
	Blue	Red
f	10 M	-1 000 M
p_f	-0.1M	-0.1M
$\alpha f + \eta$	10 M $\times \alpha + \eta$	-1 000 M $\times \alpha + \eta$
$\alpha p + \eta$	- 0.1 M $\times \alpha + \eta$	- 0.1 M $\times \alpha + \eta$
g	10 000	-1 M
p_0	0	0

Preferences incompatible with the maxmin EU model: $f \prec p_0$ and $g \succ p_0$.
Preferences incompatible with the smooth model: above, $f \sim p_f$ and $\alpha f + \eta \sim \alpha p_f + \eta$ whenever $-1000M \times \alpha + \eta < -10M$

appear to be coherent with his other behaviour: decisions where he could lose more than $\$10M$ are considered to involve very high stakes, and he is not confident enough in his advisers’ estimates to rely on them in such decisions. However, as shown in Appendix V.B, whenever the decision maker is ambiguity averse, the set of preferences just described (see Figure V.1) is inconsistent with the smooth ambiguity model. To the extent that these preferences do not seem *prima facie* unreasonable, and may even match some basic intuitions about confidence, this suggests that the smooth ambiguity model, although it allows for some violations of C-independence, cannot fully capture the role of confidence in decision making.

As we shall see in Section V.2.2, the weakening of the C-independence axiom, in a sense compatible with the considerations made above, turns out to be the central axiomatic difference between the model proposed here and comparable models. It is to this proposal that we now turn.

V.1.3 The Proposal

The first part of the proposal consists of a representation of a decision maker’s confidence in probability judgements, that is, in statements concerning probabilities, such as ‘the probability of failure of a central element in a nuclear installation is greater than 0.0002'. For this, we use a nested family of sets of probability measures, which we call a *confidence ranking*. The sets in the nested family can be thought of as corresponding to levels of confidence – the larger sets correspond to higher levels of confidence. To each set in the confidence ranking, there corresponds a set of probability judgements, namely those which hold for all probability measures in the set. This can be thought of as the set of probability judgements which the decision maker holds with the corresponding level of confidence. So for each level of confidence, the confidence ranking delivers a set of probability judgements in which the decision maker has this much confidence. Note that the larger the confidence level in question, the larger the corresponding set of probabilities measures and thus the fewer probability judgements hold for all probability measures in this set. Hence, under this representation, for higher confidence levels, fewer probability judgements are held with those levels of confidence, as one would expect. A probability judgement that holds for some set in the confidence ranking is endorsed by the decision maker, although he may have only very little confidence in it. Conversely, for a given probability judgement, the more sets in the family which are such that the probability judgement holds for all probability measures in the set, the more confident the agent is in that judgement. Just like the probability measure in the standard Bayesian theory of decision under uncertainty, a natural reading of a confidence ranking is as a subjective element capturing the decision maker's beliefs and his confidence in them.

To illustrate, suppose that a decision maker has data to estimate the probability that a certain element in a nuclear reactor will fail under extreme conditions. Upon calculation, he finds that the confidence interval is $[0.0002, 0.0005]$ at the 95% level, and that it is $[0.00005, 0.00065]$ at the 99% level. If he forms his judgements solely on the basis of these data, one might expect him to be more confident that the probability of failure ($p(F)$) is less than 0.00065 than that the probability is greater than 0.0002. This can be represented by a confidence ranking where, for every set in the ranking such that $p(F)$ is greater than 0.0002 for all probability measures in the set, $p(F)$ is less than 0.00065 for all probability measures in the set. Such a confidence ranking is illustrated in Figure V.2: the points are probability measures, the concentric circles represent the sets of probability measures in the confidence ranking and the shaded areas are the sets of probability measures where $p(F)$

Figure V.2 – Confidence ranking and confidence in probability judgements

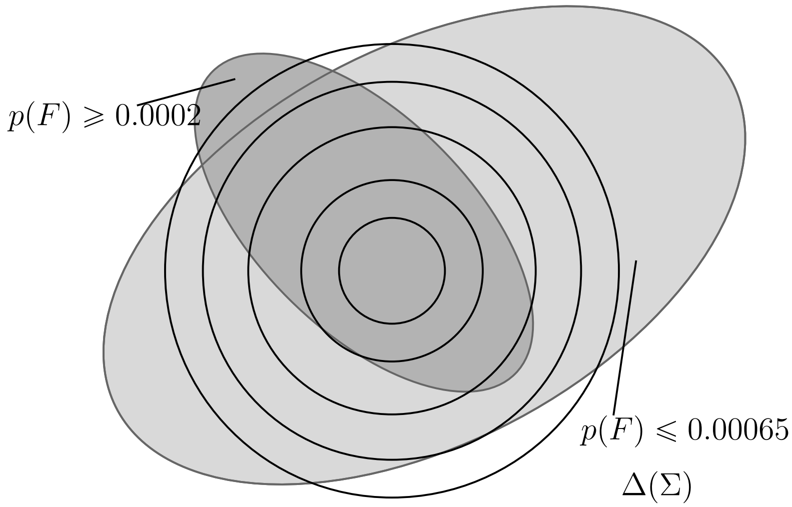

is greater than 0.0002 and less than 0.00065 respectively.[7] The fact that $p(F) \leqslant 0.00065$ holds for more of the sets in the confidence ranking that $p(F) \geqslant 0.0002$ captures the higher confidence in the former probability judgement than in the latter one.

A noteworthy property of this representation of confidence in beliefs is that it involves only an ordinal structure on the space of probability measures.[8] In this sense, the model is simpler than other recent models of decision under uncertainty, which, as discussed in Section V.1.4, require a structure on the set of probability measures with a certain degree of cardinality. Consequently, if one wished to use the model to aid decision making, one need only elicit an ordinal structure from a decision maker to capture his confidence, rather than deal with cardinal confidence comparisons.

The second part of the current proposal is a decision rule which uses this representation of confidence and which translates the maxim that the higher the stakes, the more confident one must be in a probability judgement for it to play a role in the decision. Consider a function that associates to each possible value of the stakes involved in a decision a set in the confidence ranking. We call such a function a *cautiousness coefficient*. Since a set in the confidence ranking corresponds to a level of confidence, this function can be thought of

7. It would have been more accurate to represent these latter sets of probability measures by half-planes; the current diagram is preferred because it makes the comparison between them clearer and more suggestive.

8. As shown in Proposition V.2 in Appendix V.A, a confidence ranking corresponds to an order on the space of probability measures. See also Section V.2.1.

as representing the level of confidence required for a probability judgement to play a role in a decision with given stakes. (As will be shown in Section V.3, this function captures the decision maker's attitude to taking decisions in the absence of confidence.) In particular, for any probability judgement that holds throughout the set of probability measures picked out by the cautiousness coefficient, the decision maker is confident enough in that judgement to use it in a decision with the given stakes. Hence a decision rule that uses such a function to pick out a set of probability measures and uses this set of probability measures in the decision captures the maxim mooted above.

What has been said so far constitutes the nub of the current proposal. It does not yield a single decision model but rather a recipe for generating models. In particular, two elements remain to be specified: the decision rule that the decision maker uses once he identifies the appropriate set of probability measures, and the notion of stakes he uses to assign to a decision the stakes involved in it. Different choices of decision rules and notions of stakes yield different decision models conforming to the maxim proposed above. As stated previously, the main concern here is to promote this family; however, for concreteness, we shall focus the analysis on one particular member. Namely, we concentrate on the simplest and most common decision rule using sets of probability measures, the maxmin expected utility rule (Gilboa and Schmeidler, 1989). We take as the stakes involved in the choice of an act the worst consequence that the act could yield. Both of these choices are reasonable: on the one hand, the maxmin expected utility rule has been widely vaunted as characterising careful decision making, whereas on the other hand it seems reasonable that the worse the worst consequence of an act, the more the decision maker has to lose on choosing it, and so the higher the stakes. Section V.4 contains a brief discussion of other possible decision rules and notions of stakes.

Concretely, for a cautiousness coefficient D, and a function σ assigning to each act a level of stakes that varies inversely with the worst potential consequence of choosing the act, we consider a decision rule such that the agent (weakly) prefers act g to act f if and only if:

$$\min_{p \in D(\sigma(f))} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D(\sigma(g))} \sum_{s \in S} u(g(s)).p(s) \tag{V.1}$$

To illustrate this rule, consider the decision about the sort of technology to use in a nuclear installation. To decide according to the rule, a decision maker has to first ascertain the stakes involved in the choice, which, in the specific model considered here, correspond to the worst possible consequence. For example, the stakes could correspond to a major nuclear accident. Then he needs to determine how much confidence to require in a probability judgement to use it in a decision of such gravity. If his confidence over probability measures is calculated solely on the basis of data using confidence intervals (as in the example above), this corresponds to assigning a confidence level which he will use. For example, he may decide that the appropriate set of probability measures will be those corresponding to the confidence interval at the 99.9% level. Then he would evaluate a technology by the minimum expected utility taken over this set, which, given the very low utility of a nuclear accident, boils down to using the worst-case probability for an accident in the interval.

Note that, were the decision maker considering whether to use a similar technology in another installation with extreme conditions – a smelting plant, for example – this procedure allows him to use the same data differently. For example, if the consequences of failure in the smelting plant are deemed less serious than the consequences of failure in a nuclear installation, then he might allow himself to use probability judgements in which he is less confident. Concretely, this would translate into using, say, the confidence interval at a 95% level, which may lead to different choices of technology.

It is clear that this sort of rule can capture the sort of behaviour mentioned in the example given at the beginning of the paper and developed in Section V.1.2. Since the worst consequence of allowing the planting of GM crops (or of bet f in Figure V.1) is worse than that of allowing the construction of the factory (respectively, of g), the stakes are higher, so more confidence is required, which implies the use of a larger set of probability measures. As noted in Section V.1.2, the use of a more pessimistic probability estimate, which is an immediate consequence of a larger set of probability measures in the context of the maxmin expected utility rule, is entirely consistent with the behaviour described in these examples.

V.1.4 Related literature

The proposed model can be thought of as a refinement of the maxmin expected utility model of Gilboa and Schmeidler (1989), according to which the decision maker maximises the minimum expected utility over a single, fixed set of probability measures. This model is too simple to capture the role of confidence in decision: were it to be thought of as a model of confidence in probability judgements, confidence would be an all-or-nothing affair – for any probability judgement, either the decision maker is fully confident of it or entirely unsure about it. Something more than a single set of probability measures is required to capture different levels of confidence; the model proposed here goes 'one level up' and takes a family of sets. One consequence of using a single set of probability measures is that the maxmin model cannot capture the relation between the confidence in the probability judgements used and the stakes involved in the decision, which is central here; another consequence is that it does not admit a clean separation of beliefs and attitude to ambiguity, whereas, as shall be shown in Section V.3, the model proposed here does. Finally, as noted in Section V.1.2, the maxmin expected utility representation implies that preferences satisfy C-independence, which appears incompatible with certain patterns of behaviour that are naturally related to confidence; by contrast, the proposed model only satisfies a weakened version of C-independence (see Section V.2.2), and can comfortably account for such patterns.[9]

By contrast with the maxmin model, other, more recent models in the literature do involve a notion of 'degree' on probability measures, but they are all structurally richer than the one proposed here. For example, Klibanoff et al. (2005); Nau (2006); Seo (2009); Ergin and Gul (2009) propose models that use a probability measure over the set of probability measures rather than a nested family of sets. In this sense, their models are cardinal at the second-order level, whereas ours is ordinal. As anticipated above, the structural parsimony of the current model may be seen as an advantage, in particular for prescriptive applications. For someone who wishes to decide using the rules just cited, he needs to fix upon a probability measure over the set of probability measures and a transformation function (above and beyond the utility function) that assigns to each expected utility value a real number. By contrast, all that is required under the proposed model is to pick out, given the stakes involved in the option, the set of probability measures that corresponds to the confidence level that the agent deems appropriate for those stakes.

Similar points hold for the variational preferences model of Maccheroni et al. (2006) as well as the confidence preferences model of Chateauneuf and Faro (2009). Like the second-order probability models mentioned above, these models require some cardinal structure on the space of probability measures, whereas only an ordinal structure is needed here.

9. The maxmin expected utility model is discussed here because it is the closest to representation (V.1). Many of the points hold for similar models, such as the α-maxmin model, which can be used to develop other members of the family of decision models described above. See also Section V.4.

Hence, for application in decision analysis, the model proposed here promises to be more tractable. Moreover, these models fail to separate beliefs from tastes for ambiguity, whereas the model proposed here does.

None of the models discussed above explicitly propose a dependence between the probability judgements used in a decision and the stakes involved in the choice; the role of various probability measures is always determined in some more complicated way. It is thus not surprising that the current model is not a special case of the aforementioned ones, as can be seen from the axiomatisation and the discussion in Section V.1.2. The same holds for the models proposed in the literature on robustness and in particular for the multiplier and constraint preferences models proposed by Hansen and Sargent (2001) (which can be thought of as special cases of the variational and maxmin models discussed above). Gajdos et al. (2008) propose an extension of the Gilboa and Schmeidler (1989) model that allows different sets of probability measures to be used in the evaluation of acts, but these sets are derived from an 'objective' set of probability measures (modelling information) which is given exogeneously, rather than being generated endogeneously from the stakes involved in the act, as is the case under the current proposal.

V.2 Representation

V.2.1 Preliminaries

We use the standard Anscombe-and-Aumann framework, as adapted by Fishburn (1970). Let S be a non-empty finite set of states, with Σ the algebra of all subsets of S, which are called *events*. $\Delta(\Sigma)$ is the set of probability measures on (S, Σ). Where necessary, we use the Euclidean topology on $\Delta(\Sigma)$. X is a nonempty set of outcomes; a *consequence* is a probability measure on X with finite support. $\Delta(X)$ is the set of consequences. Acts are functions from states to consequences; $\mathcal{A}$ is the set of acts. So, for an act f, and a state s, $f(s)$ is a lottery over X with finite support; for a utility function u over X, we will denote the expected utility of this lottery by $u(f(s)) = \sum_{x \in supp(f(s))} f(s)(x)u(x)$. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in \Re$, $0 \leqslant \alpha \leqslant 1$, the mixture $\alpha f + (1-\alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s, x) = \alpha f(s, x) + (1-\alpha)h(s, x)$ (Fishburn, 1970, Ch 13). We will often write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$.

We assume a preference relation $\preceq$ on $\mathcal{A}$; $\sim$ and $\prec$ are defined to be the symmetric and asymmetric components of $\preceq$. Null events and null states are defined in the usual way.[10] We use the following notation.

Definition V.1. For $f \in \mathcal{A}$, $\hat{f} \in \Delta(X)$ is any $f(s)$ that is $\preceq$-minimal over the non-null $s \in S$.

Intuitively, $\hat{f}$ is a constant act that is $\preceq$-equivalent to the worst consequence that f takes on non-null states.[11] Recall that this corresponds to one way of fleshing out the notion of the stakes involved in the choice of an act f: the worse the worst possible consequence of the act, the more the decision maker has to lose in choosing it and the higher the stakes involved in the choice. The particular decision rule considered below will take the stakes involved in the choice of an act to be the lowest utility that could result from the act after resolution of the uncertainty about the state of the world (that is, the utility of $\hat{f}$).[12] As noted in the Introduction, we consider this notion of stakes solely for the purposes of concreteness, and with no intention to suggest that it be considered the best or only reasonable notion of stakes; see Section V.4 for discussion of other possible notions. In the exposition below, we shall say that f has higher, lower or the same stakes as g if $\hat{f} \preceq \hat{g}$, $\hat{f} \succeq \hat{g}$ or $\hat{f} \sim \hat{g}$ respectively.

The following notion is central.

Definition V.2. A *confidence ranking* Ξ is a nested family of closed, convex subsets of $\Delta(\Sigma)$. A confidence ranking Ξ is *continuous* if, for every $C \in \Xi$, $C = \overline{\bigcup_{C' \subsetneq C} C'} = \bigcap_{C' \supsetneq C} C'$.[13] It is *centered* if it contains a singleton set.

As mentioned in the Introduction, confidence rankings represent decision makers' beliefs, and in particular their confidence in probability judgements. The sets in the confidence ranking can be thought of as corresponding to levels of confidence. The higher the

10. An event A is null if, for any acts f and f' differing only on A, $f \sim f'$. A state s is null if the singleton event $\{s\}$ is null.

11. Although $\hat{f}$ is not uniquely defined, it is defined up to $\sim$, which is all that is required for the axioms to have unambiguous meaning. This allows us to speak of $\hat{f}$ in the singular, which is useful for expositional purposes.

12. The non-nullness clause in Definition V.1 ensures that acts that differ only on null states are considered as having the same stakes, as seems reasonable.

13. For a set X, $\overline{X}$ is the closure of X. Note that the union of a nested family of convex sets is convex.

level of confidence in question, the larger the set: this translates the fact that one is confident of fewer probability judgements to that level of confidence. Confidence rankings are ordinal structures; indeed, Proposition V.2 in Appendix V.A shows that, to any confidence ranking there corresponds a weak order on the space of probability measures with appropriate properties and to any such order there corresponds a confidence ranking.

The convexity and closedness of the sets of probability measures in the confidence ranking are standard assumptions in decision rules involving sets of probabilities. The continuity property guarantees a continuity in one's confidence in probability judgements: it ensures, for example, that one cannot be confident at a certain level that probability of an event A is in $[0.3, 0.7]$ and then only confident that the probability is in $[0.1, 0.9]$ at the 'next' confidence level. [14]

The centredness property of the confidence ranking implies that, if the decision maker were forced to give his best estimate for the probability of any event, he could come up with a single value (and these values satisfy the laws of probability), although he may not be very confident in this value. In some situations this assumption may seem reasonable, whereas in others it may be less acceptable; in the representation result (Theorem V.1), there is an axiom that is necessary and sufficient for this property, so the property receives a behavioural characterisation.

Finally, we define a *cautiousness coefficient* for a confidence ranking Ξ to be a surjective function $D : \Re \to \Xi$ that preserves order, that is, such that for all $r, s \in \Re$, if $r \leqslant s$, then $D(r) \subseteq D(s)$. D assigns to a given level of stakes involved in the choice of an act a level of confidence that is required in probability judgements in order to use them in the evaluation of the act. This level of confidence corresponds to the appropriate set of probability measures in the confidence ranking. The fact that D is order-preserving corresponds to the intuition that the higher the stakes, the higher the confidence level required of probability judgements for them to play a role in the choice (and so the larger the relevant set of probability measures). Surjectivity of D basically attests to the behavioural nature of the confidence ranking: it implies that for each set of probability measures in the ranking, there will be a level of stakes, and hence an act, for which it is the relevant set for evaluating that act.

14. The continuity property used in the definition has been recognised as the appropriate notion of continuity for nested families of closed sets in Kuratowski (1938, pp 19-20). We thank Massimo Marinacci for this reference.

We consider a decision rule according to which the decision maker ranks acts f by the following criterion:

$$\min_{p \in D(\sigma(f))} \sum_{s \in S} u(f(s)).p(s)$$

where $\sigma(f) = -u(\hat{f})$ is the stakes involved in the choice of f. As noted in the Introduction, this translates the idea that for each level of stakes ($\sigma(f)$), the decision maker selects those probability judgements which he believes with confidence corresponding to those stakes (the judgements which hold for all probability measures in $D(\sigma(f))$) and only uses those in evaluating the act (he takes the maxmin over $D(\sigma(f))$).

V.2.2 The Representation

Consider the following axioms.

Axiom V-A1 (Weak order)**.** For all $f, g, h \in \mathcal{A}$: if $f \preceq g$ and $g \preceq h$, then $f \preceq h$; and $f \preceq g$ or $g \preceq f$.

Axiom V-A2 (S-Independence)**.** For all $f \in \mathcal{A}$, $c, d \in \Delta(X)$ and $\alpha \in (0, 1)$,

(i) if $d \succeq \hat{f}$, then $f \succeq c$ implies $f_\alpha d \succeq c_\alpha d$

(ii) if $d \preceq \hat{f}$, then $f \preceq c$ implies $f_\alpha d \preceq c_\alpha d$

Axiom V-A3 (Continuity)**.** For all $f, g, h \in \mathcal{A}$, the sets $\{\alpha \in [0, 1] | \ f_\alpha h \preceq g\}$ and $\{\alpha \in [0, 1] | \ f_\alpha h \succeq g\}$ are closed in $[0, 1]$.

Axiom V-A4 (S-Monotonicity)**.** For all $f, g \in \mathcal{A}$, $c, d \in \Delta(X)$ and $\alpha \in (0, 1]$ with $\hat{f} \sim \widehat{g_\alpha d}$ and $g_\alpha d \sim c_\alpha d$, if $f(s) \preceq g(s)$ for all $s \in S$, then $f \preceq c$, and if $f(s) \succeq g(s)$ for all $s \in S$, then $f \succeq c$.

Axiom V-A5 (S-Uncertainty Aversion)**.** For all $f, g \in \mathcal{A}$, $c, d \in \Delta(X)$ and $\alpha, \beta \in (0, 1)$ with $\hat{f} \sim \hat{g} \sim \widehat{(f_\alpha g)_\beta} d$, if $f \sim g \sim c$ then $(f_\alpha g)_\beta d \succeq c_\beta d$.

Axiom V-A6 (Non-degeneracy)**.** There exist $f, g \in \mathcal{A}$ such that $f \succ g$.

Axiom V-A7 (Centering)**.** There exists $c, d \in \Delta(X)$ with $d \succ c$ such that, for all $f, g, h \in \mathcal{A}$ with $\hat{f}, \hat{g}, \hat{h} \succeq c$ and all $\alpha \in (0, 1)$, $f \preceq g$ iff $f_\alpha h \preceq g_\alpha h$.

The axioms V-A1, V-A3 and V-A6 are standard. To motivate and explain the axioms V-A2, V-A4 and V-A5, it is helpful to compare them to the corresponding axioms for the maxmin expected utility model (Gilboa and Schmeidler, 1989), which are as follows:

Axiom V-A2' (C-Independence)**.** For all $f, g \in \mathcal{A}$, $d \in \Delta(X)$ and $\alpha \in (0,1)$, $f \preceq g$ iff $f_\alpha d \preceq g_\alpha d$.

Axiom V-A4' (Monotonicity)**.** For all $f, g \in \mathcal{A}$, if $f(s) \preceq g(s)$ for all $s \in S$, then $f \preceq g$.

Axiom V-A5' (Uncertainty Aversion)**.** For all $f, g \in \mathcal{A}$, $\alpha \in (0,1)$, if $f \sim g$ then $f_\alpha g \succeq f$.

Gilboa and Schmeidler (1989) established that V-A1, V-A3, V-A6, A2', A4' and A5' characterise maxmin expected utility preferences (which differ from representation (V.1) only in that the set of priors over which the decision maker is minimising is fixed rather than variable according to the stakes). As suggested in Section V.1.2, the central axiomatic difference between the standard maxmin model and the model considered here lies in the independence axiom.[15] Note firstly that mixing an act f with a constant act d is a way of changing the stakes involved: in particular, if d is below the worst consequence of f ($\hat{f}$), then mixing with d raises the stakes (the worst consequence of $f_\alpha d$ is worse than the worst consequence of f); conversely, if d is preferred to $\hat{f}$, then mixing with d lowers the stakes. In the light of this, C-independence (A2') basically states that the evaluation of an act is independent of the stakes involved. In particular, it implies that if $f \sim c$ for some constant act c, then $f_\alpha d \sim c_\alpha d$ for all α and d, which is tantamount to saying that if f is considered as indifferent to c, then they will be considered as indifferent even if evaluated after shifting the stakes up or down by any amount (by taking appropriate mixtures). However, the maxim which guides the proposal made here consists precisely in the idea that the evaluation of acts, and thus preferences between acts, *may change* upon changes in the stakes involved. From the point of view of this maxim, the C-independence axiom establishes an unwarranted independence of preferences from stakes, and as such must be weakened.

The S-independence (for Stakes-corrected independence) axiom V-A2 captures exactly the sort of weakening one would expect. The first clause basically says that if one *lowers*

15. Since C-independence is weaker than many of the independence axioms in the literature on decision under uncertainty, such as the standard independence axiom (Anscombe and Aumann, 1963) and comonotonic independence (Schmeidler, 1989), the discussion here also applies to the comparison with those models.

the stakes, by mixing an act with a constant act that is preferred to its worst consequence, then the act *cannot* be evaluated any *worse* than it was before (if it is preferred to a constant act, its mixture is preferred to the appropriate mixture of the constant act). The second clause deals with the other case: it says that if one *raises* the stakes, by mixing an act with a constant act that is less preferred than its worst consequence, then the act *cannot* be evaluated any *better* than it was before (if it is less preferred than a constant act, its mixture is less preferred than the appropriate mixture of the constant act). If one endorses the maxim proposed above and the caution inherent in the maxmin expected utility rule, this axiom is eminently reasonable: as stakes increase, acts may be evaluated as worse because one may no longer rely on probability judgements in which one does not have the required level of confidence, and likewise as stakes fall, acts may be considered as better, because one allows oneself to invoke probability judgements in which one did not previously have sufficient confidence.

To illustrate the difference between S-independence (V-A2) and C-independence (A2'), note that both imply that, for a risk-neutral decision maker,[16] if he weakly prefers a bet yielding a \$1000 loss if event A occurs and nothing if not to a sure \$250 loss (acts f and c in Figure V.3), then he would weakly prefer the bet yielding a \$50 loss if event A occurs and nothing if not to a sure \$12.5 loss (because the latter two acts are obtained from the former two by taking a 0.05-mix with the constant act yielding \$0). This is equivalent to saying that if he strictly prefers the sure loss to the bet where he could lose \$50, then he strictly prefers the sure loss when the amounts at stake are twenty times larger. However, C-independence also imposes the converse relationship: if the decision maker strictly prefers the sure loss when he could lose \$1000 by betting, then he strictly prefers the sure loss if he could only lose \$50 by betting. This is not necessarily reasonable, for the fact that he avoids the bet when the stakes are high (when he could lose \$1000) does not imply that he will not be willing to bet when the stakes are lower (when he has less to lose). S-independence only implies a relationship between the preferences in one direction, but not in this other, more controversial, direction.

As anticipated above, this axiom is the main difference between the proposed decision rule and many of the standard rules in the literature on decision under uncertainty. In par-

16. Alternatively, the example can be reformulated in terms of utiles instead of dollars, or using an urn whose composition is determined by whether the event occurs or not, so that it holds irrespective of risk attitudes.

Figure V.3 – S-independence

	A	A^c
f	-1000	0
c	-250	-250
g	-50	0
d	-12.50	-12.50

C-independence: $d \succ g \Leftrightarrow c \succ f$.

S-independence: $d \succ g \Rightarrow c \succ f$.

ticular, if S-independence is replaced by C-independence, the representation (V.1) reduces to the standard maxmin expected utility representation.

The S-monotonicity (for Stakes-corrected monotonicity) axiom V-A4 can be understood as a distilled version of the standard monotonicity axiom A4', extracting the aspects that are independent of stakes. The standard axiom establishes an implication from a (state-wise) dominance relation between a pair of acts to the appropriate preference over them; however, in the light of the intuitions mooted above, there are two considerations supporting it. One involves a possible difference in stakes: since g dominates f, the stakes involved in g can only be lower than those involved in f, and this might contribute to it being preferred. The other is independent of this difference: even if g had the same stakes as f, it would be preferred. S-monotonicity (V-A4) separates out the latter reason – which corresponds more closely to the standard intuition behind the monotonicity axiom – from the former. Indeed, it states, in essence, that if g dominates f then g is preferred to f when evaluated as if it had the same stakes as f, and likewise for the case where g is dominated by f. To see this, note that, since mixing acts with constant acts may shift the stakes, one can use such mixtures to glean information about preferences at different levels of stakes. For example, for an act g and a constant act c, the preferences over the mixtures $g_\alpha d$ and $c_\alpha d$ indicate the preferences between g and c at the stakes level given by the worst consequence of $g_\alpha d$. We can thus speak of preferences over g and c at a given stakes level, understanding this as shorthand for speaking of preferences over appropriate mixtures of g and c. Read in this way, the assumption in the S-monotonicity axiom states that g is indifferent to c at the stakes level corresponding to f; in other words, c can be thought of as the 'certainty equivalent' of g

when evaluated at the stakes level corresponding to f. And so the first clause of the axiom states that if g dominates f, then it is preferred to f when it is evaluated at the stakes level of f, whereas the second clause states that if g is dominated by f, then f is preferred to it when it is evaluated at the stakes level of f. Hence S-monotonicity can indeed be thought of as encapsulating the intuition behind the standard monotonicity axiom, whilst refining it to account for the issue of stakes. In the presence of the other axioms (notably V-A2), it is easily seen to be a strengthening of the standard axiom.

Similar remarks apply to the S-uncertainty aversion (for Stakes-corrected uncertainty aversion) axiom V-A5: it can be read as stating essentially the same thing as the standard uncertainty aversion axiom A5', but for a fixed stakes level. In this case, the restriction on the stakes level enters twice. On the one hand, the standard axiom applies to acts that are indifferent, whereas the acts may have different stakes (and would not necessarily be indifferent if evaluated as if they had the same stakes). S-uncertainty aversion only applies to acts f and g that have the same stakes. On the other hand, the standard axiom demands that the mixture $f_\alpha g$ be preferred to f; but, as in the case of the monotonicity axiom, such preferences could be defended by appeal to the stakes involved in the acts (the stakes involved in the mixture can only be lower than those involved in the initial acts), or by appeal to the preferences after having been 'corrected for' stakes. S-uncertainty aversion isolates the second aspect, only demanding that $f_\alpha g$ be preferred to f at the stakes level corresponding to f (where the notion of preference at a given stakes level is cashed out in terms of mixtures as explained above). Summing up these two points, the axiom basically states that whenever two acts having the same stakes are indifferent, then any mixture is weakly preferred to the two acts when evaluated as if it had their stakes level; as claimed, this is essentially the content of the standard uncertainty aversion axiom, restricted to a single stakes level. As such, V-A5 is faithful to the principal intuition behind the standard axiom, which concerns preference for hedging. Indeed, the main cases where it differs from the standard axiom involve no hedging. An example is given in Figure V.4: whilst the standard axiom implies that if a risk-neutral decision maker [17] is indifferent between acts f and g in Figure V.4, then he prefers h (which is a $\frac{3}{4}$-$\frac{1}{4}$ mixture of f and g) to both, S-uncertainty aversion is more liberal – it allows this pattern of preferences, without requiring

17. As above, the example could be reformulated in terms of utiles instead of dollars to avoid assumptions about risk attitudes.

Figure V.4 – Uncertainty aversion

	A	A^c
f	400	-100
g	50	-25
h	312.50	-81.25

Standard uncertainty aversion: $f \sim g \Rightarrow h \succeq f$.

it. As such, the axiom proposed here seems more reasonable: whilst the preference pattern seems perfectly *permissible*, it is hard to see why it should be an *obligatory* constraint on preferences, as the standard axiom demands. Certainly, given the absence of hedging considerations in this example, the normal intuition behind uncertainty aversion cannot explain why a decision maker who shows indifference between the first pair of acts *must* prefer the mixture in this case.

Like their counterparts elsewhere in the literature on decision under uncertainty, V-A2, V-A4 and V-A5 are fully behavioural, and in principle testable.

The final axiom, Centering (V-A7), states that the standard independence axiom is satisfied on a subset of acts where one cannot do worse that a certain consequence. (The condition of existence of a $d > c$ guarantees that this subset is non-trivial.) That is, it states that, for low enough stakes, the independence axiom is satisfied. As will be clear from the theorem, the axiom is not required for the representation, but if added guarantees that the confidence ranking representing the preferences is centered.

These axioms are necessary and sufficient for a representation of preferences by a decision rule featuring confidence, as represented by an confidence ranking.

Theorem V.1. *The following are equivalent:*

(i) $\preceq$ *satisfies V-A1–V-A6,*

(ii) there exists a nonconstant utility function $u : X \to \Re$*, a continuous confidence ranking* Ξ *and a cautiousness coefficient* $D : \Re \to \Xi$ *such that, for all* $f, g \in \mathcal{A}$*,* $f \preceq g$ *iff*

(V.1) $$\min_{p \in D(\sigma(f))} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D(\sigma(g))} \sum_{s \in S} u(g(s)).p(s)$$

where $\sigma(f) = -u(\hat{f})$ *for all* $f \in \mathcal{A}$*.*

Furthermore, for any other nonconstant utility function u'*, continuous confidence ranking* Ξ' *and cautiousness coefficient* D' *representing* $\preceq$ *as in (ii), there exists a positive real number* a *and a real number* b *such that* $u' = au + b$*,* $\Xi' = \Xi$ *and* $D'(x) = D(\frac{x}{a} - \frac{b}{a})$ *for all* $x \in \Re$*.*

Finally, under the conditions above, $\preceq$ *satisfies V-A7 iff* Ξ *is centred.*

Proofs of all results are in Appendix V.A.

V.3 Confidence, cautiousness and ambiguity

The proposed model involves a strict separation of beliefs and tastes. Of the three elements in the model – the utility function, the confidence ranking and the cautiousness coefficient – the first, as is standard, captures the decision maker's tastes over outcomes. As discussed above, the confidence ranking represents not only the probability judgements endorsed by the decision maker (his beliefs) but also the confidence with which he endorses them (his confidence in his beliefs). The third element, the cautiousness coefficient, captures the decision maker's attitude to choosing in the absence of confidence. The aim of this section is to demonstrate that these interpretations are indeed valid.

First, let us consider whether the confidence ranking has separate behavioural meaning from the cautiousness coefficient, by asking when two preference relations satisfying the axioms in Theorem V.1 are represented by the same utility functions and confidence rankings (though potentially different cautiousness coefficients). The following definition and proposition provide an answer.

Definition V.3. Let $\preceq^1$ and $\preceq^2$ be preference relations satisfying the axioms V-A1–V-A6. $\preceq^1$ and $\preceq^2$ are *confidence equivalent* if (i) for all $c, d \in \Delta(X)$, $c \preceq^1 d$ iff $c \preceq^2 d$ and (ii) for each $d \in \Delta(X)$, there exists $d' \in \Delta(X)$ and $\alpha \in (0,1)$ such that, for all $f \in \mathcal{A}$ and $c \in \Delta(X)$ with $\hat{f} \sim d$, $f \succeq^1 c$ iff $f_\alpha d' \succeq^2 c_\alpha d'$.

Proposition V.1. *Let $\preceq^1$ and $\preceq^2$ be preference relations satisfying the axioms V-A1–V-A6 and represented by utility functions, confidence rankings and cautiousness coefficients (u_1, Ξ_1, D_1) and (u_2, Ξ_2, D_2) respectively. $\preceq^1$ and $\preceq^2$ are confidence equivalent if and only if u_1 is a positive affine transformation of u_2 and $\Xi_1 = \Xi_2$.*

Confidence equivalence can be understood as the behavioural translation of the fact for two decision makers to have the same confidence *in preferences*. If f is preferred to c at a given stakes level but g is not preferred to d at that level, whilst it is preferred to d at a lower stakes level, then this can be interpreted as indicating that the decision maker is more confident in his preference for f over c than in his preference for g over d. In light of the interpretation of preferences at a given stakes level in terms of preferences between mixtures introduced in Section V.2.2, confidence equivalence thus states that the two decision makers share the same relative judgements of confidence in their preferences: one is more confident in the preference for f over c than in the preference for g over d if and only if the other one is. To the extent that confidence in preferences is determined by confidence in beliefs and confidence in utilities (which is trivial in this model, due to the use of a single utility function), one would expect that if two decision makers have the same confidence in preferences, then they have the same confidence in beliefs and utilities. Proposition V.1 is the formal confirmation of this intuition. [18]

Confidence equivalence guarantees that the decision makers' confidence rankings are the same, but it does not guarantee that they are used in the same way. In particular, it says nothing about which sets in the confidence ranking are used to evaluate acts with given stakes. This is exactly where their attitudes to choosing in the absence of confidence comes in. To introduce this notion, consider the standard approach to risk attitudes.

18. The notion of confidence equivalence is sufficient for the purposes of this section, where the focus is on separating attitudes towards choosing in the absence of confidence from confidence itself. Nevertheless, it is possible to behaviourally separate utility from confidence in beliefs in much the same way as, for example, Savage (1954) does. Under his theory, decision makers 1 and 2 have the same beliefs if, for any consequences c, c', d, d' such that $c \prec_1 d$ and $c' \prec_2 d'$, 1's preferences over acts whose outcomes are in $\{c, d\}$ are the same as 2's preferences over 'corresponding' acts whose outcomes are in $\{c', d'\}$. (More precisely, if f, g take values in $\{c, d\}$ and f', g' are defined by: for all $s \in S$, $f'(s) = c'$ if $f(s) = c$ and $f'(s) = d'$ if $f(s) = d$, and likewise for g', then $f' \preceq_2 g'$ iff $f \preceq_1 g$.) A similar condition, extended to consequences in $\{\alpha c + (1-\alpha)d | \alpha \in [0,1]\}$ and $\{\alpha c' + (1-\alpha)d' | \alpha \in [0,1]\}$ respectively, and relativised to stakes as in Definition V.3, ensures that decision makers have the same confidence rankings, although potentially different utilities.

In one standard treatment of risk attitude (Pratt, 1964), risk attitude is explored by considering the decision maker's attitudes to the adding of small risks to a sure amount. This sort of approach is not usually followed in the case of decision under uncertainty because of the difficulty of getting a clean notion of what it would be to add a particular 'uncertainty', especially given the dependence of the perceived uncertainty on the decision maker's beliefs. However, in the context of the current theory, where confidence enters into play via the stakes involved in the act, there is a simple analogue to adding risks: increasing the stakes.

Consider two decision makers, 1 and 2, with the same utility function and confidence ranking, and consider an act f and a constant act c. Since 1 and 2 have the same confidence rankings, at sufficiently small stakes they will both agree on the preferences between f and c; suppose that f is preferred to c at small stakes. As the stakes increase, the decision makers may switch and cease to prefer f to c; the difference between them lies in the stakes at which they cease to exhibit this preference. For example, if there is a stakes level where 1 no longer prefers f to c but 2 still does, this means that at this stakes level, 1 does not have enough confidence in his probability judgements to sustain the preference, whereas 2 does. But since they have the same confidence ranking over probability judgements, this must be because 1 is demanding more confidence in probability judgements for them to play a role at these stakes than 2. In other words, he is less comfortable than 2 in choosing with limited confidence at this stakes level; he is exhibiting more aversion to choosing in the absence of confidence than 2. If this is the case for all acts f and constant acts c, we say that 1 is more averse to choosing in the absence of confidence than 2. This leads to the following comparative notion of aversion to choosing in the absence of confidence.

Definition V.4. Suppose that $\preceq^1$ and $\preceq^2$ satisfy axioms V-A1–V-A6 and that they are confidence equivalent. Then $\preceq^1$ is *more averse to choosing in the absence of confidence* than $\preceq^2$ if, for all $f \in \mathcal{A}$, $c, d, e \in \Delta(X)$ and $\alpha \in (0, 1]$, if $f_\alpha d \succeq^1 c_\alpha d$ whenever $\widehat{f_\alpha d} \succeq^1 e$, then $f_\alpha d \succeq^2 c_\alpha d$ whenever $\widehat{f_\alpha d} \succeq^2 e$.

This definition translates formally the motivation above: it states that 1 is more averse to the absence of confidence if whenever he prefers f to c for a given level of stakes and lower, then so does 2.

As noted in Section V.1.4, the model proposed here can be thought of as a contribution to the literature on ambiguity. In this literature, discussions of comparative ambiguity aversion are generally motivated by considerations of the sets of acts that are preferred to particular constant acts (as in Yaari (1969)). This often leads to notions of comparative ambiguity aversion where, under assumptions guaranteeing that the agents have the same beliefs, a decision maker is more ambiguity averse than another if, every time the former prefers an act to a constant act, then the latter does too. The following is a definition of comparative ambiguity aversion which is true to the spirit of those notions. [19]

Definition V.5. Let $\preceq^1$ and $\preceq^2$ be preference relations satisfying the axioms V-A1–V-A6 and suppose that they are confidence equivalent. Then $\preceq^1$ is *more ambiguity averse* than $\preceq^2$ if, for any $f \in \mathcal{A}$ and $c \in \Delta(X)$, if $f \succeq^1 c$ then $f \succeq^2 c$.

The two comparative notions of aversion – to choosing in the absence of confidence and to ambiguity – are equivalent. Moreover, they are characterised entirely in terms of the cautiousness coefficient.

Theorem V.2. *Suppose that $\preceq^1$ and $\preceq^2$ satisfy axioms V-A1–V-A6 and are confidence equivalent. Suppose that they represented by (u, Ξ, D_1) and (u, Ξ, D_2) respectively. The following are equivalent:*

(i) $\preceq^1$ *is more averse to choosing in the absence of confidence than* $\preceq^2$

(ii) $\preceq^1$ *is more ambiguity averse than* $\preceq^2$

(iii) $D_1(r) \supseteq D_2(r)$ *for all* $r \in \Re$.

We conclude that the cautiousness coefficient captures the attitude to confidence and ambiguity, and that analysis of such attitudes can be carried out in the proposed model.

V.4 Variants and related models

Reduced to the essentials, the current proposal consists of two features: firstly, a representation of confidence in beliefs by an ordinal structure on the space of probability measures, and secondly, the idea that in the evaluation of options the decision maker uses the

19. This literature is too rich to be discussed in detail here. Schmeidler (1989) and Epstein (1999) are early treatments of the question; Ghirardato and Marinacci (2002) is closest to that used here. The assumption about beliefs is used by those whose models admit separation of beliefs and tastes (for example, Klibanoff et al. (2005)), but not by those whose models do not (for example, Maccheroni et al. (2006) and Chateauneuf and Faro (2009)).

set of probability measures in this structure which corresponds to the stakes involved. As noted in the Introduction, this delineates not a single decision model but a family of models, which differ according to the decision rule and the notion of stakes used. The principal aim of this paper is to promote this family, rather than any particular member. In the previous sections, we have considered one member with an eye to illustrating the intuitiveness of models in the family, and the interesting and attractive properties that they may have. Let us now briefly consider some other members of the family.

On the one hand, there are decision models that use a rule other than the maxmin expected utility rule employed in representation (V.1). As is well-known, this rule yields an ambiguity-averse model of decision, and the representation (V.1) is ambiguity averse for the same reason. One can obtain an ambiguity-seeking model of the sort proposed here by replacing the maxmin expected utility rule with the maxmax expected utility rule (where the minimum in representation (V.1) is replaced by a maximum). An analysis similar to that carried out in Section V.3 would apply to such a model, yielding a relation of 'more absence-of-confidence seeking'. Other decision models are obtained by replacing the maxmin expected utility rule with the α-maxmin or Hurwicz rule, which evaluates acts by a mix of maximum and minimum, or by a generalisation of that rule with non-constant α, such as that studied by Ghirardato et al. (2004).

Such modifications in the decision rule change relatively little as regards the main points made in the previous sections. Conceptually, the same representation of confidence in beliefs (confidence ranking) is used, and the set of probabilities is determined on the basis of the stakes via the cautiousness coefficient. Accordingly, there continues to be the same strict separation of beliefs (represented by the confidence ranking) and tastes (represented, at least, by the utility function and the cautiousness coefficient [20]). Practically, in the application of such rules to real-life decisions, one still needs to determine a set of probabilities on the basis of the stakes involved in the act; the difference is in how this set is used.

It is also worth noting other potentially interesting members of the family that can be seen as special cases of the model considered in the previous sections. For example, adding the appropriate analogue of co-monotonic independence would force the sets of probabilities in the confidence ranking to correspond to convex capacities (Schmeidler, 1986). In

20. For certain decision rules, there may be additional factors on the taste side; this is the case for example for the a function in the rule proposed by Ghirardato et al. (2004).

this case, the representation (V.1) would be equivalent to a stakes-sensitive Choquet expected utility representation, where the decision maker chooses according to the Choquet expected utility rule, but with a capacity that varies with the stakes. This model is particularly interesting from an operational point of view, since standard techniques that have been developed for eliciting capacities could perhaps be applied here to elicit the confidence ranking. Another potentially interesting specification involves confidence rankings where the sets are balls of different sizes around a fixed probability measure (ie. the sets are of the form $\{q \in \Delta(\Sigma) | d(p, q) \leqslant \epsilon\}$ for different ϵ, where d is the metric on the space of probability measures). These easily-parametrised models can be thought of as confidence versions of ϵ-contamination models studied elsewhere in the literature (Gajdos et al., 2008; Kopylov, 2009).

On the other hand, other members of the family of decision models may differ in the notion of stakes used. The motivation for the notion of stakes employed here is obvious, if not beyond controversy, and there are a myriad of other notions of stakes which could be considered: for example, the best consequence of the act, the difference between the best and worst consequences of the act, the minimum (or maximum) probability that the act takes values below a certain threshold (calculated with the respect to smallest or alternatively the largest set of probability measures in the confidence ranking) or the minimum expected utility of values below the threshold (calculated with the respect to smallest or the largest set of probability measures in the confidence ranking). This is not to mention notions of stakes that are not functions of individual acts but of the menu of available acts; for example, one could take as stakes the worst consequence out of all the acts on the menu. A full discussion of all the options is beyond the scope of this paper. However, many of the points made above continue to hold for other notions of stakes. Changing the notion of stakes used makes little difference conceptually: confidence in beliefs, represented by confidence rankings, continues to play a role in choice via the cautiousness coefficient, with only the allocation of stakes to acts (the function σ in (V.1)) changing. Accordingly, the strict separation of beliefs and tastes is retained. Finally, the application of the model in practice proceeds as described above, the only difference being the properties of acts that are relevant in determining the set of probabilities used.

Let us close with two remarks. First of all, notice that, in principle, the decision rule on sets of probabilities and the notion of stakes can be varied independently. For example, any combination of maxmin expected utility and maxmax expected utility rule with stakes as worst consequence and stakes as best consequence is possible. Whilst this yields a very large family of decision models, only some will be of interest. For example, one might wonder whether there is any interest for someone who takes as stakes the best consequence to be so pessimistic as to use the maxmin expected utility rule. Secondly, notice that both of these issues are behaviourally meaningful. Both the specific decision rule used and the particular notion of stakes employed will have consequences for properties of preferences and so may, in principle, be gleaned from choices. It remains a largely open question how to axiomatise different combinations of decision rule and notion of stakes or whether there is a general axiomatic framework in which all models in the family can be treated.

V.5 Conclusion

This paper proposes a model of a decision maker's confidence in his probability judgements in terms of a confidence ranking – a nested family of sets of probability measures. This model can be used to formulate a family of decision rules which allow the probability judgements that play a role in a decision to depend on what is at stake. Under each of the decision models in this family, the decision maker evaluates acts using a set of probability measures that is determined by the stakes involved and his cautiousness coefficient. Models in this family differ according, firstly, to the way the set of probability measures is used in the evaluation and, secondly, to the notion of stakes employed. However, all share the advantages of the general framework: structural simplicity of the representation of confidence, relative tractability in application to aiding decision making, and separation of beliefs and tastes. A particular model in the family is studied in detail and found to have a reasonable axiomatisation. Finally, a notion of comparative aversion to choosing in the absence of confidence is defined, characterised in terms of the cautiousness coefficient and related to standard notions of ambiguity aversion.

V.A Proofs of results

Throughout the Appendix, B will denote the space of all real-valued functions on S, and $ba(S)$ will denote the set of real-valued set functions on S, both under the Euclidean topology. Recall that, under this topology, $ba(S)$ is locally convex (Aliprantis and Border, 2007, §5.12). B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $a \in B$, we define $\sigma(a) = -\min_{non-null s \in S} a(s)$, and for $x \in \Re$, we define x^* to be the constant function taking value x. Finally, for any $a \in B$ and $r \in \Re$, let $e_a^r \in B$ be such that $e_a^r(s) = a(s) + \sigma(a) - r$ for all $s \in S$ and let $x_a^r \in \Re$ be $r - \sigma(a)$. By definition, $\sigma(e_a^r) = r$ and $a = e_a^r + x_a^{r*}$.

Proof of Theorem V.1

The main part of the result is the sufficiency of the axioms for the representation (the direction (i) to (ii)), the proof of which proceeds as follows. As is standard, the axioms imply the existence of a utility function u on $\Delta(X)$, with image K, and a functional I on $B(K)$ such that, for any $f, g \in \mathcal{A}$, $f \preceq g$ iff $I(u \circ f) \leqslant I(u \circ g)$. The main lemma, Lemma V.A.3, establishes that for any stakes level there exists a closed and convex set of probability measures such that each act with these stakes is evaluated by the maxmin expected utility over this set. By Lemma V.A.4, these sets form a nested family, with larger stakes yielding larger sets, and hence is a confidence ranking. By Lemma V.A.5, this confidence ranking is continuous.

We now proceed with the proof of the Theorem. We assume first (i); we will show (ii). Consider firstly the following lemmas, all of which are under the assumption of axioms V-A1–V-A6.

Lemma V.A.1. *There exists a non-constant utility function u representing the restriction of $\preceq$ to the constant acts. Moreover, for $K = u(\Delta(X))$ and $B(K)$ the set of functions in B taking values in K, there exists $I : B(K) \to \Re$ such that, for any $f, g \in \mathcal{A}$, $f \preceq g$ iff $I(u \circ f) \leqslant I(u \circ g)$.*

Proof. First note that, for all $c, d \in \Delta(X)$, if $c \sim d$ then $\frac{1}{2}c + \frac{1}{2}e \sim \frac{1}{2}d + \frac{1}{2}e$ for all $e \in \Delta(X)$. If $e \succeq c \sim d$, then, by the first clause of V-A2, it follows that $\frac{1}{2}c + \frac{1}{2}e \succeq \frac{1}{2}d + \frac{1}{2}e$; moreover, by the symmetric argument, $\frac{1}{2}d + \frac{1}{2}e \succeq \frac{1}{2}c + \frac{1}{2}e$, whence $\frac{1}{2}c + \frac{1}{2}e \sim \frac{1}{2}d + \frac{1}{2}e$, as required.

If $e \preceq c \sim d$, then the second clause of V-A2 can be applied in a similar way to obtain $\frac{1}{2}c + \frac{1}{2}e \sim \frac{1}{2}d + \frac{1}{2}e$, as required. Given V-A1, V-A3 and V-A6, it follows that the axioms of the Herstein and Milnor (1953) mixture theorem are satisfied, hence the existence of the required utility function.

Moreover, by V-A1, V-A3 and V-A4, for any act f there exist constant acts c, c' such that with $c \preceq f \preceq c'$. (If $f(s) \sim \hat{f}$ for all non-null $s \in S$, then take $c = c' = \hat{f}$. Now suppose that this is not the case. Take any $c, c' \in \Delta(X)$ such that $c \preceq f(s) \preceq c'$ for all $s \in S$. By V-A3, there exists $\alpha \in (0, 1]$ such that $\hat{f} \sim c_\alpha c'$; V-A4 implies that $f \succeq c$. Moreover, if $c \prec \hat{f}$, then $\alpha \in (0, 1)$, so V-A4 implies that $f \preceq c'$. Finally, if $c \sim \hat{f}$, then applying the previous argument on $f_\beta c'$ yields that $f_\beta c' \preceq c'$ for all $\beta \in [0, 1)$; by V-A3, $f \preceq c'$.) By V-A3 and V-A1, using a standard argument, u can be extended to a function $J : \mathcal{A} \rightarrow \Re$ such that, for any $f, g \in \mathcal{A}$, $f \preceq g$ iff $J(f) \leqslant J(g)$. It also follows from V-A4 that $J(f) = J(f')$ whenever $f(s) \sim f(s')$ for all $s \in S$, so there exists a function $I : B(K) \rightarrow \Re$ such that $I(u \circ f) = J(f)$ for all $f \in \mathcal{A}$, as required. □

Without loss of generality, it will be assumed that $0 \in K$ but not on its boundary. So K is an interval on the real line containing 0 and such that 0 is not on its boundary. Let $-K = \{r \in \Re | -r \in K\}^o$.[21] Note finally that, by construction, $I(x^*) = x$.

The following lemma will prove useful.

Lemma V.A.2. *For all $a, a', b \in B$ and $\alpha > 0$, such that $a' - \sigma(b)^*, a - \sigma(b)^* \in B(K)$, $\sigma(a) = \sigma(a') = 0$ and $a = \alpha a'$, $I(a - \sigma(b)^*) + \sigma(b) = \alpha(I(a' - \sigma(b)^*) + \sigma(b))$.*

Proof. Consider the case of $\alpha \in (0, 1]$ (the case of $\alpha > 1$ follows as a consequence). Take $g \in \mathcal{A}$ such that $a' - \sigma(b)^* = u \circ g$, $x, y \in \Delta(X)$ such that $u(x) = -\sigma(b)$ and $g \sim y$, and $f \in \mathcal{A}$ such that $f = \alpha g + (1-\alpha)x$. So $u \circ f = \alpha u \circ g + (1-\alpha) u \circ x = \alpha(a' - \sigma(b)^*) - (1-\alpha)\sigma(b)^* = a - \sigma(b)^*$. Then V-A2 applies, yielding that $f = \alpha g + (1-\alpha)x \sim \alpha y + (1-\alpha)x$, whence $I(a - \sigma(b)^*) = I(u \circ f) = \alpha I(u \circ y) + (1-\alpha) I(u \circ x) = \alpha I(a' - \sigma(b)^*) + (1 - \alpha) I(-\sigma(b)^*)$, and so $I(a - \sigma(b)^*) + \sigma(b) = \alpha(I(a' - \sigma(b)^*) + \sigma(b))$ as required. □

The following is the central lemma of the proof.

Lemma V.A.3. *For every $r \in -K$, there exists a closed convex set $\mathcal{C}_r \subseteq \Delta(\Sigma)$ such that, for every $a \in B(K)$ such that $\sigma(a) = r$, $I(a) = \min_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s)$.*

21. For a set X, X^o is its interior.

Proof. For each $b \in B(K)$ with $\sigma(b) \in -K$, we construct a probability measure p_b such that $I(b) = \sum_{s \in S} b(s)p_b(s)$ and $I(a) \geqslant \sum_{s \in S} a(s)p_b(s)$ for all a with $\sigma(a) = \sigma(b)$. To this end, take any $b \in B(K)$ with $\sigma(b) \in -K$ and define $I_b : B \to \Re$ as follows:

(i) $I_b(x^*) = x$ for all $x \in \Re$;

(ii) $I_b(a) = I(a - \sigma(b)^*) + \sigma(b)$ for all $a \in B$ such that $a - \sigma(b)^* \in B(K)$ and $\sigma(a) = 0$;

(iii) $I_b(a) = \alpha I_b(a')$ for all $a \in B$ such that $\sigma(a) = 0$ and $a = \alpha a'$ with $a' - \sigma(b)^* \in B(K)$.

(iv) $I_b(a) = I_b(e^0_a) + I_b(x^{0*}_a)$ for all $a \in B$ such that $\sigma(a) \neq 0$.

Lemma V.A.2 guarantees that I_b is well-defined (and in particular, that there is no contradiction between the clauses (ii) and (iii)), and that it is homogeneous of degree one on $\{a \in B | \ \sigma(a) = 0\}$. Moreover, it is superadditive: that is, for all $a, \bar{a} \in B$, $I_b(\frac{1}{2}a + \frac{1}{2}\bar{a}) \geqslant \frac{1}{2}I_b(a) + \frac{1}{2}I_b(\bar{a})$. This can be seen as follows.

Consider firstly the case where $\sigma(a) = \sigma(\bar{a}) = 0$. If $I_b(\bar{a}) = 0$, then the desired inequality holds by V-A2, V-A4 and the definition of I_b, and similarly for $I_b(a) = 0$; so we may henceforth suppose that $I_b(a), I_b(\bar{a}) > 0$. Without loss of generality, it can assumed that $a - \sigma(b)^*, \bar{a} - \sigma(b)^*, \sigma(\beta a + (1-\beta)\bar{a})^* - \sigma(b)^* \in B(K)$; if not, multiply the non-$\sigma(b)$ terms below through by a small positive factor and use the fact that I_b is homogenous of degree one on $\{a \in B | \ \sigma(a) = 0\}$. If $I_b(a) = I_b(\bar{a})$, then, for any $\beta \in (0,1)$, $I_b(\beta a + (1 - \beta)\bar{a}) = 2I_b(\frac{1}{2}(\beta a + (1-\beta)\bar{a}) + \frac{1}{2}(\sigma(\beta a + (1-\beta)\bar{a}))^*) + I_b((-\sigma(\beta a + (1-\beta)\bar{a}))^*) = 2I\left(\frac{1}{2}(\beta(a - \sigma(b)^*) + (1-\beta)(\bar{a} - \sigma(b)^*)) + \frac{1}{2}(\sigma(\beta a + (1-\beta)\bar{a})^* - \sigma(b)^*)\right) + \sigma(b) - \sigma(\beta a + (1-\beta)\bar{a}) \geqslant 2\left(\frac{1}{2}I(a - \sigma(b)^*) + \frac{1}{2}I(\sigma(\beta a + (1-\beta)\bar{a})^* - \sigma(b)^*)\right) + \sigma(b) - \sigma(\beta a + (1-\beta)\bar{a}) = I_b(a)$, where the first, second and last equalities follow from the definition of I_b and its homogeneity of degree one on $\{a \in B | \ \sigma(a) = 0\}$, and the middle inequality results from an application of V-A5 (with acts corresponding to $a - \sigma(b)^*$, $\bar{a} - \sigma(b)^*$, $\sigma(\beta a + (1-\beta)\bar{a})^* - \sigma(b)^*$, and $I(a - \sigma(b)^*)^*$ in the place of f, g, d and c respectively). If $I_b(a) > I_b(\bar{a})$, let $\alpha = \frac{I_b(\bar{a})}{I_b(a)}$, so that $\sigma(\alpha a) = \sigma(\bar{a}) = 0$ and $I_b(\alpha a) = I_b(\bar{a})$. It follows, by the reasoning for the case where $I_b(a) = I_b(\bar{a})$, that $I_b(\beta \cdot \alpha a + (1-\beta)\bar{a}) \geqslant I_b(\alpha a) = \beta I_b(\alpha a) + (1-\beta)I_b(\bar{a}) = \beta \cdot \alpha I_b(a) + (1-\beta)I_b(\bar{a})$ for every $\beta \in (0,1)$. Substituting $\frac{1}{1+\alpha}$ for β and multiplying both sides by $\frac{1+\alpha}{2\alpha}$ yields the desired inequality. The case where $I_b(a) < I_b(\bar{a})$ is treated similarly.

Consider finally any $a, \bar{a} \in B$. Note that, by the definition of σ, $\sigma(\frac{1}{2}a + \frac{1}{2}\bar{a}) = \frac{1}{2}\sigma(a) + \frac{1}{2}\sigma(\bar{a}) + \sigma\left(\frac{1}{2}(a + \sigma(a)) + \frac{1}{2}(\bar{a} + \sigma(\bar{a}))\right)$. So $e^0_{\frac{1}{2}a+\frac{1}{2}\bar{a}} = \frac{1}{2}a + \frac{1}{2}\bar{a} + \sigma(\frac{1}{2}a + \frac{1}{2}\bar{a}) = \frac{1}{2}a + \frac{1}{2}\bar{a} +$

$\frac{1}{2}\sigma(a)+\frac{1}{2}\sigma(\bar{a})+\sigma\left(\frac{1}{2}(a+\sigma(a))+\frac{1}{2}(\bar{a}+\sigma(\bar{a}))\right)=\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}}+\sigma(\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}})=e^0_{\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}}}$. By the definition of I_b and x^0_a, we thus have that $I_b(\frac{1}{2}a+\frac{1}{2}\bar{a})=I_b(e^0_{\frac{1}{2}a+\frac{1}{2}\bar{a}})-\sigma(\frac{1}{2}a+\frac{1}{2}\bar{a})=I_b(e^0_{\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}}})-\sigma(\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}})-\frac{1}{2}\sigma(a)-\frac{1}{2}\sigma(\bar{a})=I_b(\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}})-\frac{1}{2}\sigma(a)-\frac{1}{2}\sigma(\bar{a})$. By the result established above, $I_b(\frac{1}{2}e^0_a+\frac{1}{2}e^0_{\bar{a}})-\frac{1}{2}\sigma(a)-\frac{1}{2}\sigma(\bar{a})\geqslant\frac{1}{2}I_b(e^0_a)+\frac{1}{2}I_b(e^0_{\bar{a}})-\frac{1}{2}\sigma(a)-\frac{1}{2}\sigma(\bar{a})$. So $I_b(\frac{1}{2}a+\frac{1}{2}\bar{a})\geqslant\frac{1}{2}I_b(a)+\frac{1}{2}I_b(\bar{a})$, as required.

To construct the required probability measure, there are two cases to consider.

Case i: $I(b)>I((-\sigma(b))^*)$. For $\bar{b}=b+\sigma(b)^*$, let $D^b_1=\{a\in B\mid I_b(a)>I_b(\bar{b})\}$ and $D^b_2=co(\{a\in B\mid a\leqslant I_b(\bar{b})^*\}\cup\{\bar{b}\})$.[22] Evidently, both sets are convex and nonempty (in the case of D^b_1, convexity is ensured by the superadditivity of I_b).

We now show that D^b_1 and D^b_2 are disjoint; that is, for any $a\in B$ and $\beta\in[0,1]$, if $a\leqslant I_b(\bar{b})^*$, then $\beta a+(1-\beta)\bar{b}\notin D^b_1$, ie. $I_b(\beta a+(1-\beta)\bar{b})\leqslant I_b(\bar{b})$. The result is immediate for $\beta=0,1$, so consider $\beta\in(0,1)$. Let $a'\in B$, $x\in\Re$ be such that $\beta a+(1-\beta)\bar{b}=a'+x^*$, with $\sigma(a')=0$. Without loss of generality, it can assumed that $a'-\sigma(b)^*,(1-\beta)(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*,(-\beta I_b(\bar{b})-\sigma(b)+x)^*\in B(K)$; if not, multiply the non-$\sigma(b)$ terms below through by a small positive factor and use the fact that I_b is homogenous of degree one on $\{a\in B|\ \sigma(a)=0\}$. Since $a\leqslant I_b(\bar{b})^*$, $a'+x^*\leqslant\beta I_b(\bar{b})^*+(1-\beta)\bar{b}$. Hence $a'-\sigma(b)^*\leqslant(1-\beta)(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*$. Moreover, $\frac{1}{2}\left((1-\beta)(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*\right)+\frac{1}{2}(-\beta I_b(\bar{b})-\sigma(b)+x)^*)=\frac{1-\beta}{2}\bar{b}-\sigma(b)^*$ and $\sigma(\frac{1-\beta}{2}\bar{b}-\sigma(b)^*)=\sigma(a'-\sigma(b)^*)$. By Lemma V.A.2, $I(\frac{1-\beta}{2}\bar{b}-\sigma(b)^*)=\frac{1-\beta}{2}I(\bar{b}-\sigma(b)^*)-\frac{1+\beta}{2}\sigma(b)$. Given the preceding facts, and the fact that $\frac{1-\beta}{2}I(\bar{b}-\sigma(b)^*)-\frac{1+\beta}{2}\sigma(b)=\frac{1}{2}\left((1-\beta)I(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*\right)+\frac{1}{2}(-\beta I_b(\bar{b})-\sigma(b)+x)^*)$, V-A4 can be applied (with the acts corresponding to $a'-\sigma(b)^*$, $(1-\beta)(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*$, $(1-\beta)I(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})^*$ and $(-\beta I_b(\bar{b})-\sigma(b)+x)^*$ in the place of f, g, c and d respectively) to yield the conclusion that $I(a'-\sigma(b)^*)\leqslant(1-\beta)I(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})$. Since $(1-\beta)I(\bar{b}-\sigma(b)^*)+\beta(I_b(\bar{b})-\sigma(b)-\frac{x}{\beta})=I_b(\bar{b})-\sigma(b)-x$, it follows that $I_b(a')+x\leqslant I_b(\bar{b})$; so $I_b(a'+x^*)\leqslant I_b(\bar{b})$, as required.

Given the aforementioned properties, a separation theorem (Aliprantis and Border, 2007, Theorem 7.30) implies that there is a nonzero linear functional on B, P_b, and a real number α such that $P_b(c)\leqslant\alpha\leqslant P_b(a)$ for all $a\in D^b_1$ and $c\in D^b_2$.

$I_b(\bar{b})$ is strictly positive, by the choice of b and the definition of I_b; it follows by the

22. For a set X, $co(X)$ is the convex hull of X.

definition of D_1^b and D_2^b that α is strictly positive. Without loss of generality, it can be assumed that $\alpha = I_b(\bar{b})$. Since $I_b(\bar{b})^* \in D_2^b$, $P_b((I_b(\bar{b}))^*) \leqslant I_b(\bar{b})$; but since $(I_b(\bar{b}))^*$ is a limit point of D_1^b, $P_b((I_b(\bar{b}))^*) \geqslant I_b(\bar{b})$ and so $P_b((I_b(\bar{b}))^*) = I_b(\bar{b})$. Hence, by linearity of P_b, $P_b(1^*) = 1$. Moreover, for all $E \in \Sigma$, $(I_b(\bar{b}))_{E^c}0 \leqslant (I_b(\bar{b}))^*$ (where $x_E y$ is the function taking value x on E and y elsewhere) since $I_b(\bar{b})$ is positive; so $(I_b(\bar{b}))_{E^c}0 \in D_2^b$ and hence $P_b((I_b(\bar{b}))_{E^c}0) \leqslant I_b(b)$. Since $P_b((I_b(\bar{b}))_E 0) + P_b((I_b(\bar{b}))_{E^c}0) = P_b((I_b(\bar{b}))^*) = I_b(\bar{b})$, it follows that $P_b((I_b(b))_E 0) \geqslant 0$. Hence, by the linearity of P_b, that $P_b(1_E 0) \geqslant 0$. By a classic duality result, there is a probability measure p_b such that $P_b(a) = \sum_{s \in S} a(s)p_b(s)$ for all $a \in B$.

We now show that, for all $a \in B(K)$ such that $\sigma(a) = \sigma(b)$, $P_b(a) \geqslant I(a)$, with equality for $a = b$. For any such a, and for all $x > I_b(\bar{b}) - I_b(a + \sigma(b)^*)$, $I_b(a + \sigma(b)^*) + I_b(x^*) > I_b(\bar{b})$, and so $I_b((a + \sigma(b)^*) + x^*) > I_b(\bar{b})$ by the definition of I_b; hence $(a + \sigma(b)^*) + x^* \in D_1^b$. By construction of P_b, we thus have $P_b((a + \sigma(b)^*) + x^*) \geqslant I_b(\bar{b})$; by linearity of P_b, it follows that $P_b(a + \sigma(b)^*) + P_b(x^*) \geqslant I_b(\bar{b})$. By the continuity of P_b (letting $x \to I_b(\bar{b}) - I_b(a + \sigma(b)^*)$), it follows that $P_b(a + \sigma(b)^*) \geqslant I_b(a + \sigma(b)^*)$. By the linearity of P_b and definition of I_b, this implies that $P_b(a) \geqslant I(a)$. Since $\bar{b} \in D_2^b$, $P_b(\bar{b}) \leqslant I_b(\bar{b})$, so $P_b(\bar{b}) = I_b(\bar{b})$ and hence $P_b(b) = I(b)$ as required.

Case ii: $I(b) \leqslant I((-\sigma(b))^*)$. Let $\bar{b} = b + \sigma(b)$, and take $D_1^b = \{a \in B \mid I_b(a) > 0\}$ and $D_2^b = \{\beta\bar{b} \mid \beta \in [0,1]\}$. Evidently, both sets are convex and nonempty (in the case of D_1^b, convexity is ensured by the superadditivity of I_b). Note furthermore that, by V-A4, $I(b) = I((-\sigma(b))^*)$, so $I_b(\bar{b}) = 0$; moreover, it follows from Lemma V.A.2 that $I_b(c) = 0$ for all $c \in D_2^b$, so D_1^b and D_2^b are disjoint.

By a separation theorem (Aliprantis and Border, 2007, Theorem 7.30), there is a nonzero linear functional on B, P_b and a real number α such that $P_b(\bar{b}) \leqslant \alpha \leqslant P_b(a)$ for all $a \in \overline{D_1^b}$. Since $0^* \in D_2^b$ and it is a limit point of D_1^b, α must be 0.

For all $E \in \Sigma$, $0^* \leqslant 1_E 0$ (where $x_E y$ is the function taking value x on E and y elsewhere), so, by V-A4 and V-A2, $0 \leqslant I_b(1_E 0)$, hence $P_b(1_E 0) \geqslant 0$. By a classic duality result, there is a probability measure p_b such that $P_b(a) = \sum_{s \in S} a(s)p_b(s)$ for all $a \in B$.

Using a similar argument to that used in case (i), it can be shown that, for all $a \in B(K)$ such that $\sigma(a) = \sigma(b)$, $P_b(a) \geqslant I(a)$, with equality for $a = b$.

For every $r \in -K$, define $\mathcal{C}_r$ to be the closure of the convex hull of $\{p_b \mid \sigma(b) = r\}$.

By construction, $I(a) = \min_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s)$ for any $a \in B(K)$ such that $\sigma(a) = r$, as required.

□

The following two lemmas guarantee that the sets of probability measures mentioned in Lemma V.A.3 have the properties needed to constitute a continuous confidence ranking.

Lemma V.A.4. *For all $r, s \in -K$, $\mathcal{C}_r \subseteq \mathcal{C}_s$ if $r \leqslant s$.*

Proof. For any $a \in B(K)$ and for any $r, s \in -K$ such that $r \leqslant s$, there exists $x \in K$ and $\alpha \in (0, 1]$ such that $\sigma(\alpha(e^s_a) + (1-\alpha)x^*) = r$. Since $x \geqslant -s$, V-A2 implies that, for any $y \in K$, if $I(e^s_a) \geqslant y$, then $I(\alpha(e^s_a) + (1-\alpha)x^*) \geqslant \alpha y + (1-\alpha)x$. By the linearity of the maxmin functional, it follows that $\min_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s) = \frac{1}{\alpha} \min_{p \in \mathcal{C}_r} \sum_{s \in S} (\alpha(e^s_a) + (1-\alpha)x^*)(s)p(s) + (x^s_a - \frac{1-\alpha}{\alpha}x) \geqslant \frac{1}{\alpha}(\alpha \min_{p \in \mathcal{C}_s} \sum_{s \in S}(e^s_a)(s)p(s) + (1-\alpha)x) + (x^s_a - \frac{1-\alpha}{\alpha}x) = \min_{p \in \mathcal{C}_s} \sum_{s \in S} a(s)p(s)$. Similarly, by considering a multiple of the reflexion of a in 0^* ($a' \in B(K)$ such that, for some α, $a(s) = -\alpha a'(s)$ for every $s \in S$), we have that $\max_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s) \leqslant \max_{p \in \mathcal{C}_s} \sum_{s \in S} a(s)p(s)$. So $[\min_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s), \max_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s)] \subseteq [\min_{p \in \mathcal{C}_s} \sum_{s \in S} a(s)p(s), \max_{p \in \mathcal{C}_s} \sum_{s \in S} a(s)p(s)]$ for all $a \in B(K)$. It follows, by Ghirardato et al. (2004, Proposition A.1), that $\mathcal{C}_r \subseteq \mathcal{C}_s$, as required. □

Lemma V.A.5. *For any $r \in -K$, $\mathcal{C}_r = \bigcap_{r' > r} \mathcal{C}_{r'} = \overline{\bigcup_{r' < r} \mathcal{C}_{r'}}$.*

Proof. We show that $\mathcal{C}_r = \bigcap_{r' > r} \mathcal{C}_{r'}$; the proof that $\mathcal{C}_r = \overline{\bigcup_{r' < r} \mathcal{C}_{r'}}$ is similar. By Lemma V.A.4, $\mathcal{C}_r \subseteq \mathcal{C}_{r'}$ for all $r' > r$. Suppose, for reductio, that $\mathcal{C}_r \subsetneq \bigcap_{r' > r} \mathcal{C}_{r'}$ for some $r \in -K$, so that there exists a point (probability measure) $p \notin \mathcal{C}_r$, but $p \in \bigcap_{r' > r} \mathcal{C}_{r'}$. By a separation theorem (Aliprantis and Border, 2007, 5.80), there is a nonzero (continuous) linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) \leqslant \alpha < \phi(q)$ for all $q \in \mathcal{C}_r$. Since S is finite (so B is finite-dimensional), B is reflexive, and, by the standard isomorphism between $ba(S)$ and B^*, it follows that $ba(S)^*$ is isometrically isomorphic to B (Dunford and Schwartz, 1958, IV.3); hence there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ, a and α can be chosen so that $\alpha \in K$, $a \in B(K)$ and $\sigma(a) = r$. Taking $f \in \mathcal{A}$ such that $u \circ f = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have, by construction, that $c \prec f$ but for any $d \in \Delta(X)$ with $d \prec c$ and $\sigma(d) > r$ and any $\alpha \in (0, 1)$, $f_\alpha d \preceq c_\alpha d \prec c$, contradicting V-A3.

□

Conclusion of the proof of Theorem V.1. To conclude the direction (i) to (ii), note that, by Lemma V.A.1, there exists a utility function u representing preferences over constant acts, and a functional I representing $\preceq$ as specified. Consider the case where K is bounded above and below; the other cases are treated similarly. Lemma V.A.3 guarantees that there exists sets $\mathcal{C}_r$ for all $r \in -K$, such that $I(a) = \min_{p \in \mathcal{C}_r} \sum_{s \in S} a(s)p(s)$ for all $a \in B(K)$ such that $\sigma(a) = r$. Define $\Xi = \{\mathcal{C}_r | \ r \in -K\} \cup \{\bigcap_{r \in -K} \mathcal{C}_r\} \cup \{\overline{\bigcup_{r \in -K} \mathcal{C}_r}\}$. Since $\mathcal{C}_r$ is closed and convex for all $r \in -K$, and since, by Lemma V.A.4, the sets $\mathcal{C}_r$ are nested, Ξ is a confidence ranking. By Lemma V.A.5, Ξ is continuous. D is defined as follows: for $r \in -K$, $D(r) = \mathcal{C}_r$, for $r < s$ for all $s \in -K$, $D(r) = \bigcap_{s \in -K} \mathcal{C}_s$ and for $r > s$ for all $s \in -K$, $D(r) = \overline{\bigcup_{s \in -K} \mathcal{C}_s}$. Order preservation and surjectivity of D are immediate from the definition and Lemma V.A.4.

Finally, for any $c \in \Delta(X)$ having the property described in the centering axiom V-A7, the restriction of $\preceq$ to $\{f \in \mathcal{A} \mid \hat{f} \succeq c\}$ is non-trivial and satisfies independence. By standard results, in the presence of independence, this implies that $\mathcal{C}_{-u(c)}$ is a singleton. So V-A7 implies that Ξ contains a singleton, as required.

The direction from (ii) to (i) is generally straightforward. The only interesting cases are S-monotonicity (V-A4) and continuity (V-A3). As regards S-monotonicity (V-A4), let $f, g \in \mathcal{A}$, $c, d \in \Delta(X)$ and $\alpha \in (0, 1]$ satisfy the conditions of the axiom. Since $\hat{f} \sim \widehat{g_\alpha d}$, $D(\sigma(f)) = D(\sigma(g_\alpha d))$; let us call this set $\mathcal{C}$. Since $g_\alpha d \sim c_\alpha d$, it follows from the representation (V.1) that $\min_{p \in \mathcal{C}} \sum_{s \in S} u(g_\alpha d(s))p(s) = u(c_\alpha d)$ and hence that $\min_{p \in \mathcal{C}} \sum_{s \in S} u(g(s))p(s) = u(c)$. If $f(s) \preceq g(s)$ for all $s \in S$, then $\min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s) \leqslant \min_{p \in \mathcal{C}} \sum_{s \in S} u(g(s))p(s)$; so $\min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s) \leqslant u(c)$ and $f \preceq c$ as required. A similar argument establishes the case where $f(s) \succeq g(s)$ for all $s \in S$. We now consider continuity (V-A3). Take any $f, g, h \in \mathcal{A}$ and consider the set $\{\alpha \in [0, 1] | \ f_\alpha h \preceq g\}$; the other case is shown similarly. Suppose that α^* is a limit point of this set, and consider a sequence (α_i) of elements in the set with $\alpha_i \to \alpha^*$. If there is a subsequence of (α_i) tending to α^* such that $\widehat{f_{\alpha_{i_n}} h} \succeq \widehat{f_{\alpha^*} h}$ for all α_{i_n} in this subsequence, then the result follows from the fact that D is order-preserving and the continuity of the maxmin EU representation (with a fixed set of probability measures). Suppose now there is no such sequence. In this case, there exists $c \in \Delta(X)$, $c \prec \widehat{f_{\alpha^*} h}$; moreover, for any such c, there exists an interval I containing α^* with $\widehat{f_\beta h} \succeq c$ for any $\beta \in I$. Hence, for each such c, there exists a subsequence $(\alpha^c_{j_n})$ of (α_i), tending to α^*, with $\widehat{f_{\alpha^c_{j_n}} h} \succeq c$ for all $n \in \mathbb{N}$. By the fact that

D is order-preserving and the continuity of the maxmin EU representation, it follows that $\min_{p \in D(-u(c))} \sum_{s \in S} u(f_{\alpha *} h(s)).p(s) \leqslant \min_{p \in D(-u(\hat{g}))} \sum_{s \in S} u(g(s)).p(s)$. Since this holds for every $c \prec \widehat{f_{\alpha *} h}$, and since, by the continuity of the confidence ranking and the surjectivity of D, $D(-u(\widehat{f_{\alpha *} h})) = \bigcap_{c \prec \widehat{f_{\alpha *} h}} D(-u(c))$, it follows that $f_{\alpha *} h \preceq g$ as required.

Finally consider the uniqueness clause. Uniqueness of u follows from the Herstein-Milnor theorem. As regards uniqueness of Ξ, proceed by reductio; suppose that u, Ξ_1, D_1 and u, Ξ_2, D_2 both represent $\preceq$ according (V.1), with $\Xi_1 \neq \Xi_2$. By surjectivity of the cautiousness coefficient, for some r, $D_1(r) \neq D_2(r)$. Suppose, without loss of generality, that $p \in D_1(r) \backslash D_2(r)$. By a separation theorem (Aliprantis and Border, 2007, 5.80), there is a nonzero linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) \leqslant \alpha < \phi(q)$ for all $q \in D_2(r)$. As in the proof of Lemma V.A.5, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ, a and α can be chosen so that $\alpha \in K$, $a \in B(K)$ and $\sigma(a) = r$. Taking $f \in \mathcal{A}$ such that $u \circ f = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have that $\min_{p \in D_1(\sigma(f))} \sum_{s \in S} u(f(s))p(s) \leqslant \min_{p \in D_1(\sigma(f))} \sum_{s \in S} u(c)p(s)$, whereas $\min_{p \in D_2(\sigma(f))} \sum_{s \in S} u(f(s))p(s) > \min_{p \in D_2(\sigma(f))} \sum_{s \in S} u(c)p(s)$, contradicting the assumption that both u, Ξ_1, D_1 and u, Ξ_2, D_2 represent $\preceq$. A similar argument establishes the uniqueness of D on $-K$; outside of $-K$, the value of D is determined (to be the smallest or largest set in Ξ), since D is order-preserving.

□

Proofs of other results

Proposition V.2. *For Ξ a (continuous) confidence ranking, $\leqslant_\Xi = \{(p,q) \in \Delta(\Sigma)^2 |$ for all $C \in \Xi, q \in C$ implies $p \in C\}$ is a quasi-convex, lower semi-continuous (and locally non-satiated except at the extremes) weak order on $\Delta(\Sigma)$.* [23]

Likewise, for $\leqslant$ a quasi-convex, lower semi-continuous (and locally non-satiated except at the extremes) weak order on $\Delta(\Sigma)$, $\Xi_\leqslant = \{\{q | q \leqslant p\} | p \in \Delta(\Sigma)\}$ is a (continuous) confidence ranking.

23. An order $\leqslant$ is quasi-convex if the lower contour sets ($\{q \in \Delta(\Sigma) | q \leqslant p\}$) are convex for all $p \in \Delta(\Sigma)$; it is lower semi-continuous if the lower contour sets $\{q \in \Delta(\Sigma) | q \leqslant p\}$ are closed for all $p \in \Delta(\Sigma)$; and it is locally non-satiated except at the extremes if, for any $p \in \Delta(\Sigma)$ such that there exists $p' > p$ and for every neighbourhood around p, there exists q in the neighbourhood such that $q > p$.

Proof. Straightforward.

□

Proof of Proposition V.1. The right to left direction is evident. For the left to right direction, the clause concerning utilities follows immediately from clause (i) of Definition V.3. The clause concerning confidence rankings follows from clause (ii) of the definition, in the light of the construction and uniqueness result in the proof of Theorem V.1. □

Proof of Theorem V.2. Let the assumptions of the theorem be satisfied.

(i) implies (iii). Suppose that $D_1(r) \not\supseteq D_2(r)$ for some $r \in (-u(\Delta(X)))^o$, and that $p \in D_2(r)\backslash D_1(r)$. Since $D_1(r)$ is a closed convex set, by a separating hyperplane theorem (Aliprantis and Border, 2007, 5.80) and as in the proof of Lemma V.A.5, there is an $f \in \mathcal{A}$ and a $c \in \Delta(X)$ such that $-u(\hat{f}) = r$, $\min_{p \in D_2(r)} \sum_{s \in S} u(f(s)).p(s) < u(c)$ and $\min_{p \in D_1(r)} \sum_{s \in S} u(f(s)).p(s) \geqslant u(c)$. So for all $f_\alpha d$ with $\widehat{f_\alpha d} \succeq \hat{f}$ $f_\alpha d \succeq^1 c_\alpha d$ whereas this is not the case for $\preceq^2$ (taking $\alpha = 1$), contradicting (i); hence (iii) is established for $r \in (-u(\Delta(X)))^o$. It follows from the order-preserving and surjectivity properties of D_1 and D_2 that (iii) holds on $r \notin (-u(\Delta(X)))^o$.

(ii) implies (i). Straightforward.

(iii) implies (ii). Consider $f \in \mathcal{A}$, with $u(\hat{f}) = r$. Since $D_1(r) \supseteq D_2(r)$, $\min_{p \in D_1(r)} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D_2(r)} \sum_{s \in S} u(f(s)).p(s)$, from which it follows that if $f \succeq_1 c$, then $f \succeq_2 c$, as required.

□

V.B The example in Section V.1.2 and the smooth ambiguity model

The smooth ambiguity model (Klibanoff et al., 2005) evaluates acts by the following functional:

$$V(f) = \int_{\Delta(\Sigma)} \phi \left(\int_S u(g(s)) d\pi(s) \right) d\mu \tag{V.2}$$

where u is a von Neumann-Morgenstern utility function, μ is a countably additive probability measure on $\Delta(\Sigma)$, and ϕ is a continuous, strictly increasing real-valued function on the image of u, which is sometimes called the transformation function. As Klibanoff et al. (2005) show, under appropriate assumptions (and with an appropriate notion of ambiguity aversion), a smooth ambiguity decision maker is ambiguity averse if and only if his ϕ is concave. Note that, if ϕ is linear, then representation (V.2) reduces to standard expected utility, and hence satisfies the C-independence axiom.

To establish the claim in Section V.1.2 that the preferences in Figure V.1 are inconsistent with (V.2) if the decision maker is ambiguity averse, we use Proposition V.3 below. This proposition states that, under certain conditions, an ambiguity averse smooth ambiguity decision maker's transformation function ϕ must be linear on a particular interval. It follows in particular that he satisfies C-independence on that interval.

Proposition V.3. *Let $\preceq$ be represented according to (V.2) with a continuous utility function u, a continuous, strictly increasing and concave transformation function ϕ and countably additive probability measure μ. Suppose that, for some non-constant $f \in \mathcal{A}$ and $b, \underline{c}, \overline{c} \in \Delta(X)$, we have that $f \sim b$ iff $f_\alpha c \sim b_\alpha c$ for all $\alpha \in (0,1]$ and $\underline{c} \preceq c \preceq \overline{c}$. Let $\underline{x} = \inf\{x |\ \exists \Pi \subseteq \Delta(\Sigma)$ with $\mu(\Pi) > 0$ and such that $\forall \pi \in \Pi,\ \int_S u(f(s))d\pi(s) \leqslant x\}$ and $\overline{x} = \sup\{x |\ \exists \Pi \subseteq \Delta(\Sigma)$ with $\mu(\Pi) > 0$ and such that $\forall \pi \in \Pi,\ \int_S u(f(s))d\pi(s) \geqslant x\}$. Then ϕ is linear on $(\min\{\underline{x}, u(\underline{c})\}, \max\{\overline{x}, u(\overline{c})\})$.*

Proof. For any act g, define the real-valued measurable function on $\Delta(\Sigma)$, $\tilde{g}$, by $\tilde{g}(\pi) = \int_S u(g(s))d\pi(s)$ (this is close to the definition of second-order act given in Klibanoff et al. (2005)); with slight abuse of notation, for a constant act b, we also use $\tilde{b}$ to denote $u(b(s))$. Noting that $\widetilde{\alpha f + (1-\alpha)c} = \alpha\tilde{f} + (1-\alpha)\tilde{c}$, it follows from representation (V.2) and the assumptions in the proposition that $\int_{\Delta(\Sigma)} \phi(\alpha\tilde{f}(\pi) + (1-\alpha)\tilde{c})d\mu(\pi) = \phi(\alpha\tilde{b} + (1-\alpha)\tilde{c})$ for all $c \in \Delta(X)$ with $\underline{c} \preceq c \preceq \overline{c}$ and all $\alpha \in (0,1]$. We first show that $\tilde{b} = \int_{\Delta(\Sigma)} \tilde{f}(\pi)d\mu$. As noted in Hardy et al. (1934, §3.18), the derivative of a continuous convex function exists except at most enumerably many points. There thus exists $\alpha \in (0,1]$ and $\underline{c} \preceq c \preceq \overline{c}$ such that $\underline{c} \preceq b_\alpha c \preceq \overline{c}$ and such that the derivative of ϕ exists at $\alpha\tilde{b} + (1-\alpha)\tilde{c}$. By the assumption in the proposition, for such α and c we have that

$$\int_{\Delta(\Sigma)} \phi(\beta(\alpha\tilde{f}(\pi) + (1-\alpha)\tilde{c}) + (1-\beta)(\alpha\tilde{b} + (1-\alpha)\tilde{c}))d\mu(\pi) = \phi(\alpha\tilde{b} + (1-\alpha)\tilde{c})$$

for every $\beta \in [0,1]$. It follows that

$$\int_{\Delta(\Sigma)} \frac{\phi\left(\beta(\alpha\tilde{f}(\pi)+(1-\alpha)\tilde{c})+(1-\beta)(\alpha\tilde{b}+(1-\alpha)\tilde{c})\right)-\phi\left(\alpha\tilde{b}+(1-\alpha)\tilde{c}\right)}{\beta\alpha(\tilde{f}(\pi)-\tilde{b})}.(\tilde{f}(\pi)-\tilde{b})d\mu=0$$

and so, taking the limit as $\beta \to 0$,

$$\int_{\Delta(\Sigma)} \frac{d\phi}{dx}(\alpha\tilde{b}+(1-\alpha)\tilde{c}).(\tilde{f}(\pi)-\tilde{b})d\mu=0$$

Where the derivative exists in the light of the remarks above. Since ϕ is strictly increasing, it follows that $\tilde{b} = \int_{\Delta(\Sigma)} \tilde{f}(\pi)d\mu$, as required. Therefore $\gamma\tilde{b} + (1-\gamma)\tilde{d} = \int_{\Delta(\Sigma)} \gamma\tilde{f}(\pi) + (1-\gamma)\tilde{d}\, d\mu$ for all $\gamma \in [0,1]$ and all $d \in \Delta(X)$.

We thus have that $\int_{\Delta(\Sigma)} \phi(\tilde{f}(\pi))d\mu = \phi(\tilde{b}) = \phi\left(\int_{\Delta(\Sigma)} \alpha\tilde{f}(\pi)d\mu\right)$. Hence McShane (1937, Theorem 5) can be applied to $\tilde{f}$, the set of second-order acts (in the sense of Klibanoff et al. (2005)), the integral with respect to μ and the function ϕ to conclude that, for all $\pi \in \Sigma(\Delta)$ except at most those belonging to a set $\mathcal{S}$ such that $\mu(\mathcal{S}) = 0$, $\tilde{f}(\pi)$ belongs to an interval I on which ϕ is linear. It follows from the definition of $\underline{x}$ and $\overline{x}$ that ϕ is linear on $(\underline{x}, \overline{x})$. Moreover, by the same argument applied to $f_\alpha c$ and $b_\alpha c$ for $\alpha \in (0,1)$ and $\underline{c} \preceq c \preceq \overline{c}$, we have that ϕ is linear on $(\alpha\underline{x} + (1-\alpha)u(c), \alpha\overline{x} + (1-\alpha)u(c))$. Since these intervals overlap, for appropriate choices of α and c, we have that ϕ is linear on $(\min\{\underline{x}, u(\underline{c})\}, \max\{\overline{x}, u(\overline{c})\})$, as required.

□

We use this proposition to establish the claim in Section V.1.2 as follows. (Recall that, for ease, we are assuming that u is linear.) As a point of notation, let s_1 be the state where the ball drawn is blue and s_2 the state where it is red (other notation is taken from Section V.1.2; see in particular Figure V.1). First of all, since $g \sim d$ for some constant act d yielding $\$d$ with $d > 0$,[24] it follows from representation (V.2) that μ must give non-zero weight to the set of probability measures π for which $\int_S u(g(s))d\pi(s) \geqslant u(d)$. Since $d > \$0$, there exists $\alpha > 0$ and η such that $(\alpha g + \eta)(s_2) < -\$10M$ and $\alpha d + \eta > \$15000$. (It suffices to take $\eta = 0$ and α sufficiently large.) It follows from the properties established above that $\mu(\{\pi \in \Delta(\Sigma) \mid \int_S u((\alpha g + \eta)(s))d\pi(s) > u(\$15000)\}) > 0$. We shall now apply Proposition V.3 on $\alpha g + \eta$; note that the fact just established implies that $\overline{x} \geqslant u(\$15000)$

24. Recall that we are using the same symbol for the constant act and the consequence it yields. Giving that utility is linear, we further abuse of notation we use d to also denote the monetary value yielded by the constant act d.

(where $\overline{x}$ is as defined in the statement of the proposition, for the act $\alpha g + \eta$). Since $\alpha g + \eta = (10^{-3} \times \alpha) f + \eta$, the preferences specified in Section V.1.2 (see Figure V.1), and notably the fact that $\alpha' f + \eta' \sim \alpha' p_f + \eta'$ whenever $(\alpha' f + \eta')(s_2) \prec -\$10M$, imply that $(\alpha g + \eta)_\beta c \sim (\alpha(-\$100) + \eta)_\beta c$ whenever c is a constant act with $c \prec -\$10M$. Hence the assumptions of Proposition V.3 are satisfied for $\alpha g + \eta$ with any $\underline{c} \prec \overline{c} \prec -\$10M$. It follows from the Proposition that ϕ is linear on $[u(-\$1000M), u(\$10000)]$. So representation (V.2) reduces to expected utility – and in particular satisfies constant linearity – for acts taking values in that interval. The act $10^{-3} \times f + (-\$10M)$ takes values in the interval just described; moreover, since $(10^{-3} \times f + (-\$10M))(s_2) \prec -\$10M$, it follows from the preferences specified in Section V.1.2 that $10^{-3} \times f + (-\$10M) \sim -\$10000100$. Since g also only takes values in the interval described, and $g = 10^{-3} \times f$, it follows from constant linearity of the representation in this interval that $g \sim -\$100$, contradicting $g \succ \$0$. So the preferences specified are indeed incompatible with representation (V.2), under the assumption that the decision maker is ambiguity averse, as claimed.

Bibliography

Aliprantis, C. D. and Border, K. C. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition.

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Chateauneuf, A. and Faro, J. H. (2009). Ambiguity through confidence functions. *J. Math. Econ.*, 45:535–558.

Dunford, N. and Schwartz, J. T. (1958). *Linear Operators. Part I: General Theory*. Interscience Publishers, New York.

Epstein, L. G. (1999). A Definition of Uncertainty Aversion. *Rev. Econ. Stud.*, 66(3):579–608.

Ergin, H. and Gul, F. (2009). A theory of subjective compound lotteries. *J. Econ. Theory*, 144(3):899–929.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Gajdos, T., Hayashi, T., Tallon, J.-M., and Vergnaud, J.-C. (2008). Attitude toward imprecise information. *J. Econ. Theory*, 140(1):27–65.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2):133–173.

Ghirardato, P. and Marinacci, M. (2002). Ambiguity Made Precise: A Comparative Foundation. *J. Econ. Theory*, 102(2):251–289.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Hansen, L. P. and Sargent, T. J. (2001). Robust Control and Model Uncertainty. *The American Economic Review*, 91(2):60–66.

Hardy, G., Littlewood, J. E., and Polya, G. (1934). *Inequalities*. Cambridge University Press, Cambridge.

Herstein, I. N. and Milnor, J. (1953). An Axiomatic Approach to Measurable Utility. *Econometrica*, 21(2):291–297.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6):1849–1892.

Kopylov, I. (2009). Choice deferral and ambiguity aversion. *Theoretical Economics*, 4(2):199–225.

Kuratowski, C. (1938). Sur les familles monotones d'ensembles fermés et leurs applications à la théorie des espaces connexes. *Fundamenta Mathematicae*, 30:17–33.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 74(6):1447–1498.

McShane, E. J. (1937). Jensen's Inequality. *Bulletion of the American Mathematical Society*, 43:521–527.

Nau, R. F. (2006). Uncertainty Aversion with Second-Order Utilities and Probabilities. *Management Science*, 52:136–145.

Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica*, 32:122–136.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Schmeidler, D. (1986). Integral Representation Without Additivity. *Proceedings of the American Mathematical Society*, 97(2):255–261.

Schmeidler, D. (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3):571–587.

Seo, K. (2009). Ambiguity and Second-Order Belief. *Econometrica*, 77(1963):1575–1605.

Yaari, M. E. (1969). Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory*, 1(3):315–329.

VI Confidence in preferences

Abstract

Indeterminate preferences have long been a tricky subject for choice theory. One reason for which preferences may be less than fully determinate is the lack of confidence in one's preferences. In this paper, a representation of confidence in preferences is proposed. It is used to develop and axiomatise an account of the role of confidence in choice which rests on the following intuition: the more important the decision to be taken, the more confidence is required in the preferences needed to take it. This theory provides a natural account of when an agent should defer a decision; namely, when the importance of the decision exceeds his confidence in the relevant preferences. Possible applications of the notion of confidence in preferences to social choice are briefly explored. [1]

Keywords: Incomplete preferences; indeterminacy of preference; confidence in preferences; deferral of decisions; importance of decisions; choice-theoretic axiomatisations; social choice

JEL classification: D01, D71.

1. This is a modified version of a paper of same name that appeared in Social Choice and Welfare, 39(2): 273-302, July 2012. The main modifications concern Theorems VI.1 and VI.2, which incorrectly omitted the $sing$* (respectively $sing$) axioms in the original version.

Under the standard economic model, a rational agent's preferences can be represented by a complete order on the alternatives, but this has been famously and repeatedly challenged. Preferences may be fuzzy, imprecise or vague (Aumann, 1962; Salles, 1998). Preferences may be incomplete because the agent has not yet settled on the preferences which he deems appropriate, perhaps due to unresolved conflict (Levi, 1986; Morton, 1991). Or still, preferences may be incomplete because the agent does not see that some options will ever be comparable: after all, there is no reason to always expect them to be (Sen, 1997). From both a descriptive or a normative point of view, the assumption of completeness or determinacy of preferences is highly questionable.

We consider here the case of choice under certainty; the agent will be assumed to know the consequences of choosing each of the alternatives, and there will be no question of beliefs or probabilities over "states". The only relevant attitude is the agent's preferences (which, as standard, are taken to be subjective). If an agent settles on a preference for one alternative over another or decides on determinate indifference between the alternatives, we will say that he has emitted a *value assessment*: an assessment of the relative value of the alternatives for him.[2] In situations of choice under certainty, the agent's choices are standardly taken to be guided entirely by his preferences, or, to put the same point in other terms, by his value assessments. Conversely, his preferences are traditionally taken to be derivable from his choices.

Many of the challenges to the standard model mentioned above relate to the fact that agents do not always endorse clear, categorical value assessments on every pair of alternatives. One intuitive reason for this, which has been hardly emphasised though tacitly invoked at times in the literature, is that people often have differing degrees of *confidence* in their value assessments. Sometimes, they are not sure which of the alternatives is best (by their own lights). Consider moral dilemmas: an agent might be confident that he would prefer to sacrifice the life of one to save the lives of a hundred than not to; although he thinks that he would prefer to sacrifice the life of one to save the lives of five others than not to, he may be less confident in this value assessment; finally, he may be totally unsure about whether it is preferable to sacrifice the life of a gifted musician for that of a talented economist or not. The goal of this paper is to get a grip on the intuitive notion of *confidence*

2. Although we use the term 'value', we in no way intend to break with the tradition in the economic literature of considering preferences to be entirely subjective; thus the qualification that we are considering only the value of options for the agent and by his lights.

in one's preferences.

We first propose a representation of confidence in preferences (Section VI.1.1) and an account of its role in choice (Section VI.1.2). Although they may turn out to be descriptively valid, the focus is normative: assuming that it is rational to have different levels of confidence in one's preferences, the goal is to say something about what sorts of confidence one can allow oneself to have and on the role confidence should play in choice. In Section VI.2, an axiomatisation of the notion of choice on the basis of confidence in preferences is developed. Under the proposal, confidence is related to two aspects of choice situations, which though apparently relevant in many cases, have received little attention in the choice-theoretic literature to date: the importance of the decision to be taken, and the question of when and whether to defer the decision. In Section VI.3, we discuss these two issues in detail, as well as the relationship with existing literature in economics and philosophy. In Section VI.4, we turn to the possible application of the notion of confidence in social choice, attempting a preliminary investigation into the question, and proposing a social choice rule which takes into account voters' confidence in their preferences. Proofs are relegated to the Appendix.

VI.1 Preference, confidence and choice

VI.1.1 Representing confidence in preferences

Let X be a finite set of alternatives, with at least three members. Henceforth, we use the generic terms x, y and so on to refer to elements of X, and the generic terms S, T and so on to refer to subsets of X. A weak ordering on a set is a complete, reflexive, transitive binary relation on that set. The standard model represents an agent's preferences by a weak ordering on the set of alternatives X. Let $\mathcal{P}$ be the set of weak orderings on X; we use the generic terms R, R_i, and so on to refer to elements of $\mathcal{P}$ and the generic term $\mathcal{R}$ to refer to subsets of $\mathcal{P}$. The generated strict ordering and indifference relation are defined as standard.

Weak orderings represent determinate preferences: for each pair of alternatives, either the agent strictly prefers one to the other or is determinately indifferent. The most common way of representing preferences that are not determinate in this way is by weakening the completeness assumption (Sen, 1970, 1997). Reflexive, transitive relations which do not necessarily satisfy completeness are called quasi-orderings. If Q is a quasi-ordering, then there may be alternatives x and y such that neither xQy nor yQx; these are cases where the agent does not have any determinate preference – including determinate indifference – between the alternatives x and y. In other words, he does not endorse any value assessment concerning the comparison between x and y.

This is however not the only way to represent an agent who does not have determinate preferences over all pairs of alternatives. Another possibility is to use *sets of weak orderings*.[3] For a set of weak orderings $\mathcal{R}$, there may be alternatives z and w such that zRw for all $R \in \mathcal{R}$; in this case, the agent has a determinate weak preference for z over w. By contrast, there may be alternatives x and y such that neither xRy for all $R \in \mathcal{R}$ nor yRx for all $R \in \mathcal{R}$. This represents an agent who does not have any determinate preference over x and y; he endorses no value assessment concerning the comparison between these alternatives.

The representation by sets of weak orderings is strictly more expressive than the representation by a quasi-ordering in the following sense: for each set of weak orderings there is a unique quasi-ordering which represents the same preferences, but there are generally several sets of weak orderings which correspond to a given quasi-ordering. As regards the first point, given a set of weak orderings $\mathcal{R}$, define the quasi-ordering Q as follows: for all alternatives x, y, xQy if and only if xRy for all $R \in \mathcal{R}$. It is straightforward to see that Q is a quasi-ordering and that Q and $\mathcal{R}$ represent the same preferences: the agent has weak preference, strict preference, indifference or indeterminacy according to one if and only if he does according to the other. By contrast, Figure VI.1 shows two different sets of weak orderings, both of which correspond to the empty quasi-ordering (for all x, y, neither xQy nor yQx); this illustrates the fact that there may be no unique set of weak orderings corresponding to a given quasi-ordering. One can regain uniqueness by adding a constraint on the set of orderings. We say that a set of weak orderings $\mathcal{R}$ is *full* if, for any weak ordering R, $R \in \mathcal{R}$ if, for all alternatives x and y, if $xR'y$ for all $R' \in \mathcal{R}$, then xRy. This condition basically says that, if an ordering R agrees with what all orderings in $\mathcal{R}$ have in common, then R is in $\mathcal{R}$. It can be shown that to each quasi-ordering Q one can associate a unique full set of weak orderings, namely, the set containing all weak orderings R such that xRy if xQy, for all alternatives x and y.[4]

3. This method is related to that used by Sen (1973) and Levi (1986).
4. Donaldson and Weymark (1998) show that the intersection of this set of weak orderings is the initial

Figure VI.1 – Two sets of ordering corresponding to the same quasi-ordering

a	d		a	b	d	c
b	c		b	a	c	d
c	b		c	d	b	a
d	a		d	c	a	b

Both of these representations can be interpreted as representations of the agent's confidence in his preferences. He is confident in his preference for x over y if xQy, or if xRy for all $R \in \mathcal{R}$. And he has no preference concerning x and y in which he is confident if neither xQy nor yQx, or it is not the case that xRy for some $R \in \mathcal{R}$ and it is not the case that yRx for another $R \in \mathcal{R}$. As a representation of the agent's confidence in his preferences, these proposals have an evident defect: they are binary. Either the agent is completely confident in a value assessment concerning two alternatives, or he is completely unsure about any value assessment concerning them.

In reality, it seems that one can, rationally, have different degrees of confidence in one's preferences or value assessments. Take the example of moral dilemmas. An agent may be pretty confident that he prefers to sacrifice the life of one to save the lives of a thousand than to let the thousand perish. He also thinks that he prefers to sacrifice the life of one to save the lives of ten than not to, but he is less confident in this value assessment. And he is more confident in that assessment than in the following assessment which he still, perhaps cautiously, endorses: that he prefers to sacrifice the life of one "ordinary" person for the lives of ten petty criminals than not to. There thus appear to be degrees of confidence in one's value assessments or preferences; a model of confidence in preferences should be able to account for this.

This can be done by a simple extension of the second representation presented above:

---

quasi-ordering Q. Note that the "expressivity" of the notion of quasi-ordering is more appropriate for choice theory, since the information given by a choice function (under an appropriate axiomatisation) is only sufficient to pick out a unique quasi-ordering. To pick out a unique set of weak orderings, more "Boolean" information about preferences is required (for example: a is preferred to b *if* a is preferred to c). To stay closer to the traditional framework of choice theory, throughout this paper we work with the expressiveness corresponding to quasi-orderings; accordingly, everything done with sets of weak orderings will be unique only up to fullness of the sets. See also Section VI.3.3.

instead of representing preferences by a set of weak orderings, use a *nested family* of weak orderings.[5] Let Ξ be such a nested family of subsets of $\mathcal{P}$. Ξ represents confidence in preferences in the following way. If there is a set of weak orderings $\mathcal{R} \in \Xi$ such that xRy for all $R \in \mathcal{R}$, then the agent (weakly) prefers x to y. But he may not be very confident in this value assessment: his confidence in the assessment is captured by the size of the biggest set $\mathcal{R}'$ in Ξ such that xRy for all $R \in \mathcal{R}'$. So he is at least as confident in his preference for z over w than in his preference for x over y if for every set of weak orderings $\mathcal{R} \in \Xi$ such that xRy for all $R \in \mathcal{R}$, zRw for all $R \in \mathcal{R}$; and he is more confident in the former preference if there is a set $\mathcal{R}' \in \Xi$ such that zRw for all $R \in \mathcal{R}'$ but there are some $R' \in \mathcal{R}'$ for which it is not the case that $xR'y$.

Figure VI.2 illustrates the idea diagrammatically. The plane is the set of weak orderings: the points are weak orderings, so for each point and for every pair of alternatives, the alternatives are ordered one way or another according to the weak ordering corresponding to that point. The (filled) circles represent the sets in the nested family of sets representing confidence in preference; the fact that a value assessment holds in a circle means that it holds for all points (weak orderings) in that circle. Finally, the fact that a value assessment holds in a bigger circle than another represents the fact that the agent is more confident in the former than in the latter.

It is evident from the diagram that to any nested family of sets of weak orderings in $\mathcal{P}$ there corresponds a unique weak ordering *on* the set $\mathcal{P}$ of weak orderings on X. For weak orderings R and R' on X, R is lower than R' according to the weak ordering on $\mathcal{P}$ if the smallest set in the nested family containing R' contains the smallest set in the nested family containing R.[6] Intuitively, the order represents how *implausible* the weak orderings are as candidates for the "right" notion of preference (by the agent's lights): the higher a weak ordering is on the order, the "farther out" it is on the diagram in Figure VI.2, and the less the agent feels that he has to consider it as an appropriate reflection of his preferences. Implausibility is a sort of dual notion to confidence: the agent is more confident in

5. That is, a set of sets of weak orderings such that, for each pair of distinct sets, one is strictly contained in the other.

6. Formally: for Ξ a nested family of subsets of $\mathcal{P}$, define $\leqslant$ as follows: for any $R, R' \in \mathcal{P}$, $R \leqslant R'$ if, for all $\mathcal{R} \in \Xi$, if $R' \in \mathcal{R}$ then $R \in \mathcal{R}$. And for any weak order $\leqslant$ on $\mathcal{P}$, define the nested family of subsets Ξ to be that family containing all and only $\{R' | \, R' \leqslant R\}$ for all $R \in \mathcal{R}$. It is straightforward to see that this is a bijection from nested families of subsets of $\mathcal{P}$ to weak orderings on $\mathcal{P}$.

Figure VI.2 – Implausibility on the set of orderings

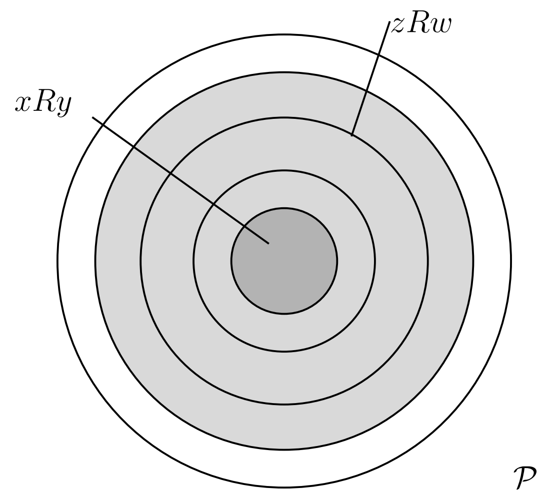

a value assessment if it holds for all weak orderings up to a higher level of implausibility, and conversely, a highly implausible weak ordering will only be taken into account if the agent demands a high level of confidence. This leads to the following representation of confidence in preferences.

Definition VI.1. An *implausibility order* $\leqslant$ is a weak ordering on $\mathcal{P}$. $\Xi_{\leqslant} = \{\{R' | \; R' \leqslant R\} | R \in \mathcal{P}\}$ is the nested family of subsets of $\mathcal{P}$ associated with $\leqslant$.

The implausibility order $\leqslant$ is said to be *centred* if there exists a single element R with $R \leqslant R'$ for all $R' \in \mathcal{P}$. This element is called the *centre*.

Henceforth, we use $\mathcal{I}$ to denote the set of implausibility orders on $\mathcal{P}$.

The rest of this paper will develop a theory of choice based on this representation of confidence in preferences. Note that the representation does impose some non-trivial conditions on the concept. In particular, it implies that for a given level of confidence, the preferences in which the agent is at least that confident are transitive and reflexive (this follows from the points made above). This is reasonable: if one is confident to a certain degree in one's preference for x over y, and one is confident to that degree in one's preference for y over z, then one is confident to at least that degree that x is preferred to z.

Centred implausibility orders have a single weak ordering as the least implausible ordering on the set of alternatives. (Equivalently, the nested family of sets contains a singleton set.) This represents the agent as having a "best guess" as to which value assessment is "right" for any pair of alternatives, though he may have very little confidence in this assessment in many cases (as represented by the rest of the implausibility order). We do not wish to take any specific position on whether this is a reasonable normative constraint on rational agents, or on whether it is descriptively accurate. The centering property of implausibility orderings will not be a requirement for most of the results presented here.

Finally, by analogy with the property of fullness of sets of weak orderings we say that an implausibility order $\leqslant$ is *full* if all the sets in $\Xi_{\leqslant}$ are full. [7]

VI.1.2 Confidence and choice

A representation of the agent's confidence in his preferences is of little use on its own; an account of the role of confidence in choice is also required. In this section, we outline the principal ideas and notions involved in this account; in Section VI.2, we axiomatise the notion of rationalisability proposed here, and in Section VI.3 we discuss in more detail some of the central notions, as well as the comparison with related literature.

The basic intuition is simple: the more important the decision to be taken, the more confident one should be in the value assessments required to take that decision. If a choice between x and y is to be made, but the choice is not particularly important, one can choose x on the basis that, on one's appraisal, x is better than y, even though one is not very confident in this value assessment. But if the choice is very important, then one needs to be a lot surer of the value assessments underlying one's decision to take it, or certainly to take it responsibly. This intuition is intended to be normative – it is intended to say something about how people should decide on the basis of value assessments in which they may be more or less confident – although a full defence is beyond the scope of this paper. It may also describe the way that people actually do make decisions in several cases, though experimental work would be required to determine to what extent this is indeed the case.

To formalise this intuition, a first requirement is a notion of the importance of a choice. We thus assume that there exists a set I of possible importance levels, and that this set is equipped with a linear ordering (that is, an antisymmetric weak ordering) $\leqslant$: $i \leqslant j$ means that the importance level j is "higher" than the level i.

The importance levels are related to two factors in a choice problem. On the one hand,

7. This can be formulated just in terms of the order itself as follows: $\leqslant$ is *full* if, for any $R, R' \in \mathcal{P}$, if $\bigcap_{R_i \leqslant R'} R_i \subseteq R$, then $R \leqslant R'$.

they are related to the degree of confidence required in a value assessment for it to play a role in the choice, via the maxim that the more important a decision, the more confident one needs to be in a value assessment for it to play a role in the choice. So to each level of importance can be associated the value assessments in which the agent has enough confidence to use for choices of this importance. Since, as discussed above, a set of such value assessments can be represented by the appropriate set of weak orderings, the relationship between importance level and confidence can be naturally represented by a function which associates to each importance level a set in the nested family of sets $\Xi_{\leqslant}$. Moreover, when the importance rises, the required degree of confidence rises, so the set of value assessments in which there is sufficient confidence becomes smaller; in the representation, this corresponds to the fact that the set of weak orderings corresponding to a higher importance level contains the set corresponding to a lower importance level. Technically, this can be captured by a function $D : I \to \wp(\mathcal{P})$ such that (i) for all $i \in I$ and all $R, R' \in \mathcal{P}$ with $R' \leqslant R$, if $R \in D(i)$, then $R' \in D(i)$, and (ii) $D(i) \subseteq D(j)$ if $i \preccurlyeq j$.

Such a function captures the agent's attitude to choosing in the absence of confidence: for two agents with the same implausibility order but different D, the one with smaller $D(i)$ requires less confidence in a value assessment to use it in a decision of importance level i than the agent with higher $D(i)$. This is a subjective factor, the agent's taste for choosing in important decisions on the basis of limited confidence, or, to put it in another way, his cautiousness when it comes to choosing on the basis of value assessments in which he has limited confidence. The function is called the *cautiousness coefficient*.

On the other hand, importance levels are supposed to capture an aspect of the choice situation or decision the agent is faced with. Some decisions are more important than others; to the former are associated importance levels that are higher (according to the order $\preccurlyeq$) than the importance levels associated to the latter. So, to each choice situation will be associated not only a set of available alternatives (sometimes called the *menu*) but also an importance level. The pair (S, i) represents the choice offered among the elements in S, with importance i. We return to this representation of choice situations and the notion of importance level in Section VI.3.1.

This only leaves the definition of choice functions. Under the standard definition, a choice function c is a function from the set of non-empty subsets of X (which we denote by $\wp(X)\backslash\varnothing$) to the set of subsets of X (denoted $\wp(X)$) such that (i) for every non-empty

$S \subseteq X$, $c(S) \subseteq S$; and (ii) for every non-empty $S \subseteq X$, $c(S)$ is non-empty. According to the maxim proposed above, an agent should chose based on value assessments which he is confident enough in given the importance of the decision; this implies that there may be decisions of such importance that he does not have sufficient confidence in the relevant value assessments to make a choice. We thus weaken the second condition and allow the choice function to yield empty choice sets. We define a *choice* function* to be a function $c : \wp(X)\backslash\varnothing \to \wp(X)$ such that $c(S) \subseteq S$ for every non-empty $S \subseteq X$. $c(S)$ is called the *choice set*, and if $x \in c(S)$ then x is said to be *admissible*. For a detailed consideration and defence of this notion of choice function, see Section VI.3.2.

The object of study are variants of choice* functions which account for importance. An *importance-indexed choice* function* is a function $c : (\wp(X)\backslash\varnothing) \times I \to \wp(X)$ such that $c(S, i) \subseteq S$ for every non-empty $S \subseteq X$ and every $i \in I$.

Having introduced this new sort of choice function, a corresponding notion of rationalisability is required. The idea is simple: for each choice situation, the importance level picks out, via the cautiousness coefficient D, a set of weak orderings which represent all the value assessments in which the agent is confident enough to use in his choice. He chooses on the basis of this set of weak orderings in a specified way. We thus propose a notion of rationalisability of a choice* function by a set of weak orderings, which is then extended to a notion of rationalisability of an importance-indexed choice* function by an implausibility order.

Definition VI.2. For any $S \in X$ and $\mathcal{R} \subseteq \mathcal{P}$, let $\sup(S, \mathcal{R}) = \{x \in S|\ xRy \text{ for all } y \in S \text{ and all } R \in \mathcal{R}\}$.

A choice* function c is *rationalisable by a set of weak orderings* if there exists $\mathcal{R} \subseteq \mathcal{P}$ such that, for all non-empty $S \subseteq X$, $c(S) = \sup(S, \mathcal{R})$.

An importance-indexed choice* function c is *rationalisable by an implausibility order* if and only if there exists an implausibility order $\leqslant$ and a cautiousness coefficient D such that, for all non-empty $S \subseteq X$ and $i \in I$, $c(S, i) = \sup(S, D(i))$.

The set $\sup(S, \mathcal{R})$ contains those elements of S which are at least as good as all the other elements of S according to all the weak orderings in $\mathcal{R}$. Rationalisability by a set of weak orderings $\mathcal{R}$ says that an element is in the choice set if and only if it is at least as good as all other elements on the menu according to all the weak orderings in $\mathcal{R}$. Rationalisability by an implausibility order says that, for every importance level i, an element is in the choice set if it is at least as good as all the other alternatives according to all orderings in the set corresponding to that importance level, $D(i)$.

The notion of rationalisability by a set of weak orderings proposed above has received little attention in the choice-theoretic literature. Much more popular is the notion according to which the choice set contains those elements which are best according to at least one ordering, rather than according to all orderings; in other words, where the choice set is the *union* of the sets of best elements according to each of the weak orderings, rather than the *intersection* (Moulin, 1985).[8] The intersection notion proposed above is of course stronger than the union notion, but it is traditionally seen as problematic, because, unlike the union notion, it does not always yield non-empty choice sets. However, this property, though it may be unwanted if one is interested in rationalising choices by a single ordering or by a single set of orderings, is less problematic for implausibility orders. All that the emptiness of the choice set indicates is that there are degrees of confidence such that the agent is not confident of any particular choice to that degree. This does not imply that he cannot make a choice – he could always choose, but in doing so he would have to rely on preferences in which he may not be very confident. We shall return to this issue in detail in Section VI.3.2.

VI.2 Representation

In this section we give necessary and sufficient conditions for rationalisability by an implausibility order. To this end, consider the following properties of importance-indexed choice* functions c.

For all $x, y \in X$, $S, T \subseteq X$ and $i, i' \in I$,

8. Translated in terms of quasi-orderings, the notion of rationalisibility proposed here picks out the set of *optimal* elements, to use Sen's (1997) terminology, whereas the union notion picks out the set of *maximal* ones. As noted in the text, maximal elements of quasi-orderings always exist, whereas this is not the case for optimal elements.

α^*	If $x \in S \subseteq T$ and $x \in c(T,i)$, then $x \in c(S,i)$
π^*	If $x \in S$, $y \in S \cap T$, $y \in c(T,i)$ and $x \in c(S,i)$, then $x \in c(S \cup T, i)$
$sing^*$	$x \in \gamma(\{x\}, i)$
Consistency	If $x \in c(S,i)$ and $i \geqslant i'$, then $x \in c(S,i')$
Centering	There exists $j \in I$ such that $c(S,j)$ is non-empty

We have the following result.

Theorem VI.1. *An importance-indexed choice* function is rationalisable by an implausibility order if and only if it satisfies* α^**,* π^**,* $sing^*$ *and Consistency. Moreover, it is rationalisable by a centred order if and only if it satisfies Centering. In both cases, there is a unique coarsest full rationalising implausibility order and cautiousness coefficient.* [9]

The proof is to be found in the Appendix. It relies heavily on a representation result for choice* functions, which involves the following two properties.

α	if $x \in S \subseteq T$ and $x \in c(T)$, then $x \in c(S)$
π	if $x \in S$, $y \in S \cap T$, $y \in c(T)$ and $x \in c(S)$, then $x \in c(S \cup T)$
$sing$	$x \in \gamma(\{x\})$

Theorem VI.2. *A choice* function* c *is rationalisable by a set of weak orderings if and only if it satisfies* α*,* π *and* $sing$*. Moreover, in this case, there is a unique full rationalising set of weak orderings. Finally, if* c *always takes non-empty values, then the rationalising set of weak orderings is a singleton.*

Evidently, the properties α^*, π^* and $sing^*$ in the representation of importance-indexed choice* functions are just the importance-indexed versions of α, π and $sing$. They state that α, π and $sing$ hold whenever the importance level is the same. As concerns the properties themselves, the first is Sen's α (also called Chernoff's property) and requires no further discussion. The last merely demands that when there is no real choice to be made – the menu is a singleton – there is an admissible alternative – namely, the element in the menu. Of course, this property is automatically satisfied by standard choice functions. By contrast,

9. Recall that an implausibility order $\leqslant$ is coarser than $\leqslant'$ if, for any $R, R' \in \mathcal{P}$, $R \leqslant' R'$ implies that $R \leqslant R'$, but $R <' R'$ does not necessarily imply that $R < R'$. For a definition of fullness, and a discussion of its relevance here, see Section VI.1.1 and in particular footnote 4.

to our knowledge, there has been little study of choice* functions; accordingly, the property π and Theorem VI.2 are new.

To illustrate, π says that if x is a best candidate for a position from a European university and y is a best candidate from an American university, and if y is also affiliated to a European university, then x is a best candidate from among European and American universities. π^* says that this consequence holds whenever the choices all have the same importance level. It follows from the final clause in Theorem VI.2 that, on choice functions, π is equivalent to Sen's β. However, in the absence of the non-emptiness condition, π is strictly stronger than β. On the one hand, π implies β: for $x, y \in c(S)$, $S \subseteq T$ and $y \in c(T)$, π applies to x, y, $S = S \cap T$ and T, yielding that $x \in c(S \cup T) = c(T)$ as required. On the other hand, here is an example where β is satisfied but π is not: $X = \{x, y, z\}$, $c(\{x, y\}) = \{x\}$, $c(\{y, z\}) = \{y\}$, $c(\{x, z\}) = \{x\}$ and $c(\{x, y, z\}) = \{\}$. It follows from the theorem above that β is too weak to guarantee rationalisation of choice* functions by sets of weak orderings; π is the appropriate property for choice* functions.

Among the properties in Theorem VI.1, consistency is doubtless the one which differs most from the traditional ones in the literature. For good reason: it concerns the comparison between choices at different levels of importance. It says that any option which is admissible when the importance is high will continue to be admissible when the menu remains the same but the importance level drops. In other words, as the importance decreases, more alternatives become admissible – and so may be chosen – but no previously admissible alternatives cease to become admissible. Of course, as is standard in choice theory, the fact that an alternative is admissible does not mean that it will *actually* be chosen. So this property is compatible with (concrete) cases where the option actually chosen when the decision is important is not that which is chosen when it becomes less important: it only demands that the alternative could (rationally) have been chosen in the less important situation.

The final property, Centering, states that one can always make a choice from any menu, provided the importance level is low enough. In many cases, this might seem reasonable: although one is not confident enough in one's relevant value assessments to pick out an option when the decision is important, one has no trouble selecting some "best guesses" when little rests on the decision. This property is only required for the implausibility order to be centered (Theorem VI.1); as noted in Section VI.1.1, we do not wish to take any position here on the centredness of the implausibility order, and correspondingly on whether the Centering property on importance-indexed choice* functions is normatively advisable or descriptively acceptable in general.

Note finally that this representation, and Theorem VI.1, is a strict generalisation of the standard theory of choice and the axiomatisation by Sen's properties α and β. If $c(S, i) = c(S, j)$ for all importance levels i and j and all non-empty subsets S, then the four properties above equivalent to the conjunction of α and β. The cautiousness coefficient sends all the importance levels to the same, singleton, set of weak orderings, so the representation collapses into the traditional representation by a weak ordering. As the cautiousness coefficient indicates, this captures the case of an agent who is insensitive to his confidence in his preferences and to the importance of the decision.

VI.3 Discussion

In this section, we first discuss in more detail two of the less standard elements of the proposal outlined above: the notion of importance level and the permissibility of empty choice sets. Then we consider the relationship, both technical and conceptual, between the current proposal and related economic and philosophical literature.

VI.3.1 The importance level

A major element of the current proposal is the extension of the ordinary representation of a choice situation from a set of available alternatives (the menu) to a set of alternatives and an importance level. The latter is exogenous, insofar as it is not derived from the menu, but taken as given along with it.[10] This extra structure might make some readers uncomfortable.

The supplementary assumptions on which this representation of choice situations relies are as follows: (i) to each choice that the agent is faced with, one can associate a set of elements from X and an importance level from I; and (ii) any pair consisting of a subset of X and an importance level from I represents a choice which the agent could conceivably

10. In choice theory, little structure is assumed on the alternatives. If more structure is assumed, it becomes possible to define an equivalent of the importance level in terms of the set of alternatives on offer; see Hill (2013) for an example of how this may be done in the case of decision under uncertainty.

be faced with.

Both of these assumptions are just versions of assumptions that are involved in the traditional representation of choice situations as subsets of a set of alternatives X. On the one hand, this representation supposes that the element x when it belongs to the menu $\{x, y\}$ is in a relevant sense "the same" as the element x in $\{x, y, z\}$. This corresponds to the first assumption above, (i), which we call *identification*. On the other hand, the representation permits that all sets of elements of X represent choice situations in which the agent might conceivably find himself; this is the second aspect, (ii), which we call *richness*. [11] In practice, the choice of the set of alternatives X is at the modeller's discretion, and he has to find a balance between these two "structural" assumptions, which, though necessary in some form or other for every theory of choice, are often in tension. Consider, for instance, some of the examples Sen raises against the most natural notion of identification among alternatives (1993; 1997), such as the choice between taking tea and going home, and the extension by the offer of cocaine. [12] As Sen notes, one could reply to such examples by refining the set of alternatives to distinguish between the option of tea with cocaine not being on the menu and the option of tea with cocaine also being on the menu. However, this defence of identification leaves richness in a sorry state, for it demands that one can find situations in which the agent has the choice between some rather strange alternatives, such as between having tea with cocaine also being on the menu and going home with cocaine not being on the menu. [13]

In the light of this it is not necessarily unreasonable to impose extra structure on the representation of the choice situation: as we shall see below, this sometimes allows an improvement in identification whilst limiting the damage done by richness. Of course, to the extent that such extra structure may not be easily discernable in all decision situations, it may not be appropriate in all cases; however, this does not imply that they are no cases

11. Of course, only weaker versions of this are needed, but they all require at least that for any two elements there exists menus containing them both and representing a conceivable choice situation, and this is all that is needed for the points made below to be relevant.

12. When offered the choice between taking tea with an acquaintance and going home, the agent chooses the tea, whereas when the choice is between tea with the acquaintance, cocaine with the acquaintance and going home, he chooses home; these choices violate the property α.

13. Broome (1991, Ch. 5) discusses a related but distinct worry concerning the tension between what we have called identification and the extent to which there can exist non-empty consistency constraints on preferences.

where it is a relevant, and indeed useful, compromise between identification and richness. Here are some examples where a modelling of the sort proposed above seems reasonable:

— a governing body is considering policies for encouraging recycling in the population. It seems reasonable to say that in general the "same" policies are available (for example, advertising, fines, bonuses, nudging etc.), but that the importance of the decision differs according to whether the governing body is the head of a household or an office, local government, regional government, national government or an international body.

— a young academic is to present his work to a public of peers. The occasion could be an in-house closed seminar, an open seminar, an international conference, an occasion where only people who know his work are present, an occasion where potential employers could be in the audience and so on (the academic profile of the audience is the same in all cases). It seems that the "same" options are available concerning how to present his material, but the importance differs between the different cases.

— consider a classic moral dilemma where you have the choice between killing one person, thus saving ten, or refusing to kill the one, thus sacrificing the ten. There is a sense in which this is the "same" choice as that between killing ten people or letting a hundred die, and as that between killing hundred people and letting a thousand die, and so on; [14] but the gravity of the choices differs among these cases.

In all these cases, there is certainly a sense in which the same options are available, but the importance of the choice to be taken differs. They are thus cases to which the representation of choice situations proposed above can be applied. To show that the importance level is a factor which needs to be taken into account, it suffices to establish that the admissible choices may differ depending on the importance level. This certainly seems to be true. Although the academic may try out a less standard organisation of his presentation or incorporation of material he is less sure about on a "friendly" audience, when the importance of the event is higher it would not be unreasonable to revert to an organisation and a choice of material he is more confident in. And although it may be acceptable to try new methods of encouraging recycling on a local level, on a global level one needs to be much

14. If you prefer, replace this example with the choice between killing 0.0001% (respectively, 0.001%, 0.01% and so on) of the human population, and letting 0.001% (respectively, 0.01%, 0.1% and so on) die, or any other scaling between the cases that is deemed appropriate.

more confident to adopt them. Indeed, one sometimes hears people say that, although a policy "worked" when tried out on a local level, more reflection is needed before deciding whether to apply it nationally. Such assertions seem to rely on the tacit assumption that the national decision is more important than the local one, and so requires more deliberation. In fact, there are quite a few cases where people cite the importance of a decision as a relevant factor in the choice made. To take an example from moral theory, Rawls (1971, p169) explicitly raises the question of the importance of the agreement made under the "veil of ignorance" as a point in favour of his principles of justice; he thus admits that importance (of the choice behind the veil of ignorance as opposed to a choice taken in front of it, for example) may be relevant for the choices one takes. [15]

In many of these cases, one might have the impression that the choice is the same, but that the *context* differs. To take the first example, the same decision has to be taken about recycling, but in a household, local, regional, national or global context. This intuition can be captured by modelling the context by a function (call it γ) which associates to every menu an importance level: this is the importance attached to the choice among these alternatives in this context. The choice situations will thus be represented by pairs consisting of a menu (the alternatives on offer) and a context function (the context of the choice). This representation of choice is visibly equivalent to that proposed, and a notion of rationalisability for choice functions on pairs consisting of a set of alternatives and a context can be proposed and axiomatised as above (replacing appearances of i by $\gamma(S)$).

We take examples such as those given above to indicate that the representation of choice situations proposed in Section VI.1.2 may be relevant in several cases. Nevertheless, it is worthwhile noting that the result obtained in Section VI.2 remains valid even if the choice situation is represented in the traditional way, as a set of alternatives. Were one to represent the choice situations in the examples above in the standard way, then, as already noted, one would have to revert to "finer" alternatives. A natural choice would be to replace the set of alternatives X by the set $X \times I$ of pairs consisting of an alternative and an importance level. (x, i) is the alternative of choosing x in a choice of importance i. The importance-indexed choice* function generates a choice* function which is defined on a subset of the menus generated by this set of alternatives: namely, on those menus consisting of elements with the same importance level. In that sense, Theorem VI.1 can alternatively be thought of as

15. Thanks to Thibault Gajdos for suggesting this example.

an axiomatisation of a rationalisation of a partially-defined choice* function on more or less standard menus. The menus on which the function is not defined are those with mixed importance levels: examples include the choice between choosing advertising to promote recycling on a national level (for example, for the whole of France), and using "nudging" techniques on a local level (for the city of Caen). As noted above, it is not always easy to make sense of such choices; indeed, the fact that a (fully defined) choice function on this set of alternatives requires choices to be made on such menus is an example of the problems which too fine an identification can pose in terms of the required richness.

Of course, the representation proposed in the previous sections does not require any choices to be made on such menus, and the information gained from the choices on which it is defined does not imply any particular choices on these peculiar menus. Nevertheless, if desired, it is possible to extend the notion of rationalisation proposed above to such menus: to take just one of several possibilities, one could choose those alternatives which are best for all preference orderings singled out by the highest importance level among the alternatives on the menu. [16] Representation theorems for such notions of rationalisability can be obtained, by making appropriate modifications to Theorem VI.1 above. Depending on one's view on these sorts of mixed-importance menus, one might be more or less attracted by such theorems.

Before closing the discussion of importance levels, let us make a remark concerning the assumption that the importance levels can be linearly ordered. Basically, this boils down to assuming that the order of "higher importance" ($\preccurlyeq$) is transitive and complete. Whilst transitivity is very intuitive, completeness, though a natural assumption in many situations, may not seem to be satisfied in certain cases. To take the second example given above, it may not be possible to determine whether the talk given as an invitee to a seminar in one department (where, say, the person in question intends to apply for a position) is of higher, lower or equal importance than the talk given as an invitee in another department (which the person in question also intends to apply to); that is, it might not be possible to rank one importance level relative to the other. There is a natural generalisation of Theorem VI.1 which can deal with such cases. All that is required is a relaxation of the assumption that implausibility orders are complete: that is, that every pair of weak orderings on X can be ranked according to implausibility. [17] If the order on the importance levels $\preccurlyeq$ is transitive

16. Formally: $(x, i) \in c(S)$ if and only if $x \in \sup(S, D(\sup_{(y,j)\in S} j))$.

17. Note that, under this relaxation, $\Xi_{\preccurlyeq}$ (defined as in Definition VI.1) ceases to be a nested family of sets

but not complete, then the properties α^*, π^* and Consistency are necessary and sufficient for a rationalisation of the sort given in Definition VI.2, where the implausibility order is transitive, reflexive but not necessarily complete. The other clauses of Theorem VI.1 continue to hold.

VI.3.2 Choice* functions

It has long been recognised that indifference and indeterminacy of preferences are difficult to distinguish on the basis of choice; accordingly, the problem of "deducing" preference from choice is particularly thorny in cases where preferences may be indeterminate. Recently proposed solutions have involved weakening the Weak Axiom of Revealed Preference (Eliaz and Ok, 2006), looking at sequential choice (Mandler, 2009) or invoking choices over opportunity sets and supposing preference for flexibility (Danan, 2003). The method employed in this paper is different, and very simple: it employs choice* functions, thus relaxing the standard assumption that the choice set is necessarily non-empty. But how are the cases where the choice* function yields the empty set to be interpreted?

The simplest answer is that the agent refuses to make a decision. In practice, this may come out in many ways. For example, he might admit that he is not sure what to do. More interestingly, there may be cases where he can *defer* the decision to whoever would next have to take it (including, perhaps, his later self); this is what he would do when the choice set is empty. Deferral of decisions seems a natural option for identifying cases of incompleteness, indeterminacy, or lack of confidence in preferences. Certainly, there seem to be several non-trivial examples where deferral, or something like it, is an option:

— a secretary takes the responsibility of making many decisions on behalf of her boss without consulting him. However, there are decisions which she could be called upon to make but which she would not accept to make in the absence of her boss, or at the least without his confirmation that her proposed decision is suitable. This is a case where she does not actually make a choice from the options available, but "defers" the decision to her boss.

— in the English law system, a judge may state in his verdict that he found the case very difficult and would grant that the case is fit for appeal. (Under English law, a party who wishes to appeal has to ask the judge to declare the case fit for appeal

of weak orderings, but rather a family of sets of weak orderings partially ordered by set inclusion.

at the end of the hearing.) In essence, the judge is emitting a judgement on the case, as he must, but admitting that the case should conceivably be reconsidered by others; this is the closest thing to a deferral under the obligation to express a choice or judgement.

— a person is faced with a moral dilemma, and he is unsure about the correct option. He decides to delay taking the decision, in order to consult friends, advisers, mentors and so on on the moral issues involved. In a sense, he is deferring the decision to his future self.

A particularly well-developed study of questions relating to deferral in the economic literature is in the 'preference for flexibility' tradition, following on from the groundbreaking work of Kreps (1979) (see Danan (2003) for an application of related ideas to incomplete preferences). A specificity of this literature is that the objects of choice are taken to be menus of options. This formalism is rich enough to represent situations involving deferral; for example, the choice among the menus $\{x\}, \{y\}, \{z\}, \{x, y, z\}$ can be understood as the choice from a menu $\{x, y, z\}$ with an option to defer (represented by $\{x, y, z\}$). If anything, insofar as the study of deferral is concerned, the formalism is too rich: it is generally assumed for example that the agent can choose between (or has preferences over) $\{x, y\}$ and $\{z\}$, and it is difficult to interpret such a choice in terms of deferral. Indeed, the standard interpretation used in the literature is in terms of how the agent restricts or keeps open the future choices that he will subsequently be faced with; for example, how he may restrict or keep open his choice of meals for dinner by his choice of restaurant. This interpretation is distinct from an interpretation in terms of deferral: it assumes that one will get to make the choice in the future, whereas this need not be the case in general for deferral, as the first two examples above, where the agents are deferring to someone else, illustrate. Moreover, in harmony with the dominant interpretation, the models proposed in the literature tend to be entirely strategic: the agent makes his choice of which options to leave open solely on the basis of his beliefs about what he will prefer at the time when he will come to make his final choice. It is unclear to what extent such models provide a fully satisfactory treatment of deferral, for it is questionable whether considerations involving beliefs about future preferences exhaust the possible reasons for deferring. For example, one might defer the decision about the moral dilemma in the final example above, even if one expects friends and advisers to be of no help; likewise, a judge might effectively defer his decision in a difficult case, even if the case is so difficult that he does not expect a higher court to "do any better". It thus seems that a full account of deferral cannot limit itself to the agent's anticipations of what will be preferred (by his future self, or by another agent). There seems to be a need to incorporate the fact that one reason to defer makes no reference to expectations about future attitudes, but only to one's current attitudes: namely, that one is not sure which option to prefer.

It is an advantage of the proposal made in the preceding sections that it can be thought of as providing a theory of when to defer that responds precisely to this need. The empty choice set can be interpreted as indicating that the agent would like to defer, or that he would defer if possible. As a theory of deferral, it is eminently reasonable: it says that one should defer if one's confidence in the choice of any alternative does not match up to the importance of the decision. Under this account, deferral makes reference solely to one's current attitudes, and in particular, to one's confidence in one's (current) preferences. Of course, we do not mean to suggest that it is a complete theory of deferral: for that, one would probably require some combination of a theory such as this one with theories capturing strategic reasons for deferral, such as those mentioned above.

One might nevertheless complain that these considerations do not vindicate the use of choice* functions, for if deferral is seen as an option, then it should be incorporated into the menu offered to the agent. Indeed, this can be done, and yields a representation visibly similar to that proposed in Section VI.1.2.

Let us use the symbol † to represent the option of deferral; when † is present in the menu, the option of deferral is available, when it is absent, deferral is not available. The current proposal can be formulated entirely in terms of importance-indexed choice functions (ie. functions always yielding non-empty choice sets) on the set of alternatives $X \cup \{\dagger\}$ (where X is as above).

Now deferral is an alternative which has a special status with respect to the others. For one, the question of identification (see Section VI.3.1) is particularly complicated: whereas the alternatives are supposed to be defined at such a level of fineness that x chosen from menu S can be treated as the same x as that chosen from T, it is unclear whether there is any sense in which deferring when the choice is from menu S can be judiciously thought of as the "same thing" as deferring from the choice on menu T. In the face of this, one could introduce a set of different new alternatives "deferring from S", "deferring from T" and so on, with all the disadvantages in terms of richness that were discussed above. Alternatively, one could admit just one new alternative, †, but give it a distinguished role in the definition of rationalisability and in the axiomatisation. Since deferral is a special option, the axioms on choice will have to reflect some of its distinctive properties.

As regards rationalisability, the theory proposed above, under the interpretation of an empty choice set as deferral, immediately implies a notion of rationalisability for menus containing the deferral option †, namely: for all $S \subseteq X \cup \{\dagger\}$ such that $\dagger \in S$ and all $i \in I$, if $\sup(S\backslash\{\dagger\}, D(i))$ is non-empty, then $c(S, i) = \sup(S\backslash\{\dagger\}, D(i))$, and if not, then $c(S, i) = \dagger$. This renders explicit the idea that one does not defer if there are options which are optimal according to all the weak orderings in the relevant set and that one does defer (and not possibly do something else) if not. It remains to define the value of the choice function when deferral is not available. Of course, the notion of rationalisability proposed in Definition VI.2 does not deal with this case, but we have already mentioned an intuition about what one should do: choose an option that one is most confident in choosing. This yields the following definition of rationalisability of importance-indexed choice functions on sets of alternatives including an explicit deferral option.

Definition VI.3. An importance-indexed choice function c on a set of alternatives including an explicit deferral option, $X \cup \{\dagger\}$, is *rationalisable by an implausibility order* if and only if there exists an implausibility order $\leqslant$ and a cautiousness coefficient D such that, for all non-empty $S \subseteq X \cup \{\dagger\}$ and $i \in I$, and all $x \in X$,

$$\begin{array}{lll} x \in c(S,i) & \text{if} & x \in \sup(S\backslash\{\dagger\}, D(i)) \\ & \text{or} & \dagger \notin S \text{ and } x \in \sup(S, D(j)) \text{ for all } j \text{ s.t. } \sup(S, D(j)) \neq \varnothing \\ \dagger \in c(S,i) & \text{if} & \dagger \in S \text{ and } \sup(S\backslash\{\dagger\}, D(i)) = \varnothing \end{array}$$

The first clause says that x is in the choice set if either it is admissible by the lights of the previous notion of rationalisability (Definition VI.2) or deferral is not available and x is admissible by the lights of the previous notion of rationalisability for all levels of importance where the choice set yielded by that notion is non-empty. The second clause says that one chooses to defer if the option is available and no alternatives on the menu are admissible by the lights of the previous notion of rationalisability.

It should not be surprising that this notion of rationalisability can be axiomatised along similar lines to the axiomatisation proposed in Section VI.2. In fact, let the properties $\alpha^\dagger$, $\pi^\dagger$ and Consistency† be identical to the properties α^*, π^* and Consistency in Section VI.2, except that they apply to all $x, y \in X$ and $S, T \subseteq X \cup \{\dagger\}$, and consider the following new property and modification of Centering:

Deferral	If $\dagger \in c(S,i)$, then $c(S,i) \cap X = \varnothing$
Centering†	There exists $j \in I$ such that $c(S,j) \neq \{\dagger\}$

Deferral just states that if one defers, no alternative in X is admissible. Centering† states that for any menu there is an importance level for which one does not defer. These properties are necessary and sufficient for the rationalisability of importance-indexed choice functions where there is an explicit deferral option.

Theorem VI.3. *An importance-indexed choice function on a set of alternatives including an explicit deferral option is rationalisable by a centered implausibility order if and only if it satisfies $\alpha^\dagger$, $\pi^\dagger$, $sing^\dagger$, Consistency†, Centering† and Deferral. Moreover, there is a unique coarsest full rationalising implausibility order and cautiousness coefficient.*

We conclude that the interpretation of empty choice sets in terms of deferral is not only natural in many cases, but entirely consistent with the traditional choice-theoretic methodology, via the addition of a special option for deferral into the menus.

VI.3.3 Related Literature

The current proposal has significant technical and conceptual points of contact with the economic literature on incomplete, fuzzy or uncertain preferences, as well as with the philosophical literature on incomparability of value relations. In this section, we first discuss the technical relationships, before saying a few words on the conceptual issues.

Let us first consider the model proposed in Definition VI.1. Sets of weak orderings have been frequently used in the literature on incompleteness (for example, by Sen (1973)). Indeed, some of this literature is related to the literature on vagueness, and the technique of considering the intersection of sets of weak orderings parallels that, adopted by supervaluationist theory of vagueness (Fine, 1975), of considering sets of possible sharpenings of a vague predicate, and taking a sentence to be true if it holds under all of the sharpenings.

As such, sets of orderings feature prominently in Broome's (1997) analysis of the incommensurability of the betterness relation as vagueness. Technically, the current proposal can be thought of as a generalisation of the models used by these theorists, replacing the binary notions drawn from the single set of orderings with a relative notion, supplied by the order on the set of orderings.

Another major modelling paradigm for indeterminacy of preference, itself connected to an influential theory of vagueness, is that emanating from fuzzy set theory (Salles, 1998). There the essential modelling idea is to associate to each ordered pair of alternatives (x, y) a number in $[0, 1]$ (or, more generally, an element in an appropriate partially ordered set), which represents the degree to which x is preferred to y. Functions from pairs of alternatives to the interval $[0, 1]$ are called fuzzy binary relations (Basu, 1984; Salles, 1998). It turns out that the current proposal can also be thought of as a generalisation of a version of the fuzzy model. For example, one obtains a fuzzy binary relation from an implausibility order if one associates to each ordered pair for which a preference holds up to some "rank" in the implausibility order, the ratio of the "rank" to which the preference holds to the total number of "ranks" in the implausibility order, and if one associates zero to all other ordered pairs.[18] This mapping from implausibility relations to fuzzy binary relations involves an information loss, for much the same reason that the quasi-ordering defined from a set of weak orderings carries less information than the set itself (Section VI.1.1). To illustrate, consider an implausibility ordering according to which the leftmost pair of orderings in Figure VI.1 are more plausible than the other two orderings on the right, and these four are more plausible than all other orderings on the alternatives. Whereas the fuzzy binary relation defined above is trivial (all pairs of different elements are sent to zero), the implausibility order does contain non-trivial Boolean information: for example, that the agent is more confident that if a is preferred to b, then c is preferred to d than he is that if a is preferred to b, then a is preferred to c. This difference is of course related to the well-known

18. Formally, for an implausibility order $\leqslant$, let $n = |\Xi_{\leqslant}|$, and suppose that the elements in $\Xi_{\leqslant}$ are $\mathcal{R}_1, \ldots, \mathcal{R}_n$, where $\mathcal{R}_i \subset \mathcal{R}_{i+1}$ for each i. Then define the fuzzy binary relation on X, $r_{\leqslant} : X \times X \to [0, 1]$, as follows: for any $(x, y) \in X \times X$, $r_{\leqslant}((x, y)) = \frac{1}{n}.\max\{i | \, \forall R \in \mathcal{R}_i, \; xRy\}$ (where the maximum is taken to be zero if the set is empty). Note that this is not the only way to define a fuzzy binary relation from an implausibility order. Indeed, although $r_{\leqslant}$ is not connected, it is not difficult to define fuzzy binary relations from the implausibility order which are: an example is $r'_{\leqslant}$, defined by $r'_{\leqslant}(x, y) = \frac{1}{2}(1 + r_{\leqslant}(x, y) - r_{\leqslant}(y, x))$. It is straightforward to check that, whilst $r_{\leqslant}$ satisfies max-min transitivity (Salles, 1998), $r'_{\leqslant}$ does not.

penumbral connections in the vagueness literature (Fine (1975), see Piggins and Salles (2007) for a brief presentation): whereas under the fuzzy theory, it is not true (with degree one), for a blob situated in a vague zone between red and orange on a colour scale, that 'if the blob is not red, then it is orange', this is true under the supervaluationist theory. This relationship is natural, given the aforementioned similarities between the model proposed here and supervaluationism.

In fact, the notion of implausibility order can be seen as a cure to the ills sometimes attributed to the supervaluationist and fuzzy approaches. The fuzzy approach is particularly criticised for missing the penumbral connections (for example, Williamson (1994)); implausibility orders do not have this problem. On the other hand, some complain that the supervaluationist theory cannot do justice to the intuition that there are degrees of truth, or, in the case of models of preference, degrees of preference (Broome (1997) makes this point, and Basu (1984) makes a similar point in criticism of Sen's (1973) use of the intersection of sets of orderings). By contrast, the current proposal can cope with degrees.

A final representation, used in the literature on random utility (see, for example, Luce and Suppes (1965); Fishburn (1998)) as well as in literature on preference for flexibility (see Kreps (1979) and Section VI.3.2), involves probability functions over the space of weak orderings (or, more often, over the space of utility functions). Implausibility orders are neither weaker nor stronger than such probability functions, as can be seen by comparing the orderings on the set of non-empty sets of weak orderings generated by probability functions (the so-called 'qualitative probability relations') with the orderings naturally generated by implausibility orders.[19] In particular, whilst for the former, unlike the latter, any pair of non-minimal elements of the ordering must have non-empty intersection, the latter, unlike the former, satisfy an extra 'independence' property which guarantees additivity (see, for example, Savage (1954)). Nevertheless, there is a sense in which implausibility orderings are more parsimonious than probability functions, insofar as they are ordinal rather than cardinal.

These points pertain to the formal properties of implausibility orders; now let us consider the notion of rationalisability (Definition VI.2). The relationship with other results involving sets of orderings was discussed in Section VI.1.2. As concerns representations

19. Such an order, $\preccurlyeq$, can be defined as follows: for $\mathcal{S}, \mathcal{S}' \subseteq \mathcal{P}$, $\mathcal{S} \preccurlyeq \mathcal{S}'$ if, for all $\mathcal{R} \in \Xi_{\preccurlyeq}$, if $\mathcal{R} \subseteq \mathcal{S}$, then $\mathcal{R} \subseteq \mathcal{S}'$.

involving probabilities over the set of weak orderings, comparison is hindered by the fact that these representations require specific assumptions – generally, that the choice functions are probabilistic (Luce and Suppes, 1965; Fishburn, 1998), or that the objects of choice must themselves be sets of alternatives (Kreps, 1979; Danan, 2003) – that are not made here.

It remains to consider the literature on fuzzy preferences. As noted above, the notion of implausibility order contains more information than the corresponding fuzzy preference model. However, the supplementary information plays no role in the notion of rationalisability or the representation theorem (Theorem VI.1), as can be seen from the uniqueness clause. Hence this theorem can regarded as a representation theorem for a notion of rationalisability for fuzzy preferences. As such, it differs both in motivation and in content from existing proposals (see, for example, Salles (1998) and the references contained therein). As concerns motivation, the literature on rationalisability of fuzzy preferences contains many choice rules, and the emphasis is placed more on the search for choice rules which rationalise behaviour that is inconsistent with the standard non-fuzzy theory than on conceptual comparison and motivation of the rules. By contrast, the choice rule proposed here is based on an intuitive maxim about the role of confidence in choice. Moreover, the notions of importance level, and the idea of allowing there to be no choice one could make with particular levels of confidence are completely absent from the fuzzy preference literature. This is related to the main technical differences: whereas all representation theorems in the literature involve choice functions (which never yield empty sets) and many focus on connected fuzzy preference relations, choice* functions are involved here and the fuzzy relation derived from implausibility orders is not necessarily connected. It is instructive to compare the proposed representation with perhaps the closest rule in the literature on fuzzy preferences, namely the $B_{D[\alpha]}$-rule, according to which x is an admissible choice out of A if $r(x, y) \geqslant \alpha$ for all $y \in A$, where r is a reflexive, transitive and connected fuzzy preference relation (Dutta et al., 1986). Dutta et al. (1986) only characterise the case where the rule yields a choice function, and come to the conclusion that it is behaviourally equivalent to the standard choice rule with non-fuzzy preferences (this equivalence is related to the final clause in Theorem VI.2 above). Of course, this conclusion no longer holds if one allows empty choice sets, as Theorem VI.2 shows. Potential interesting directions for future research may be to explore the consequences for fuzzy preference theory of taking choice* functions seriously (potentially using the interpretation proposed above), and to consider the class of fuzzy preference relations generated by implausibility orders.

Before undertaking a conceptual comparison, let us note that, given its interesting technical properties, the basic modelling idea might find fruitful application beyond the case of preferences which has been considered in this paper.

Consider first two examples that immediately spring to mind given the preceding discussion. First of all, one could interpret the weak orderings as corresponding to different (precise) measures of inequality. Then the implausibility order can be thought of as, say, a judgement of the plausibility of these measures, and the formal model would provide an extension of Sen's (1973) theory of inequality that avoids the aforementioned criticism by Basu (1984). Or, to take another example, one could interpret the set of orderings on which the implausibility order is defined as betterness orderings, and the implausibility order itself in terms of truth (or, in supervaluationist terminology, as admissible sharpenings); in this way, implausibility orders could be seen as a refinement of Broome's (1997) theory of the vagueness of betterness relation, which allows both for penumbral connections and degrees of truth. Note that Broome argues that degrees of truth are not linearly ordered, whereas our implausibility order is; a simple extension of the model to allow the implausibility order to be incomplete (see Section VI.3.1) would cope with this case.

As a final example, consider Rabinowicz's (2008) analysis of value relations in terms of permissibility of preferences: he uses sets of orderings that are interpreted as containing the permissible preference orderings. If one accepts that permissibility may come in degrees, then a version of the current model – involving nested sets of preference orderings, interpreted as preferences permitted to a certain extent – would be the adequate extension of Rabinowicz's proposal. Note furthermore that Rabinowicz's framework is rich enough to capture the notion of parity (Chang, 2002); given the formal similarity, if implausibility orderings are interpreted in the terms of permissibility of preferences, the same could be said of them. Finally, Rabinowicz allows incomplete preferences in his set, in order to capture incomparability; an extension of the current model would be to take the implausibility order over the set of quasi-orderings (see Section VI.1.1) rather than weak orderings. This would make a difference in terms of the Boolean (or "penumbral") properties of the model – for example, although, under the model proposed above, it is always true that 'x is better than y or y is better than x', this is not the case if weak orderings are replaced by quasi-orderings. [20]

However the proposal in the preceding sections pertains not the the relation of inequality, or betterness, or value relations in general, but to preferences, understood in the standard economic sense, as subjective choice-guiding attitudes. What is the conceptual relationship between the notion of confidence in preferences and other accounts of incomparability of value relations, or of incompleteness, fuzziness or uncertainty of preference? Of course, from the strictly behavioural point of view standardly adopted in economics, all that can be said is what can be inferred from the comparison of the technical properties of the models, and in particular the representation theorems; that is, from the sort of comparison that we have just undertaken. The following discussion is thus for those of a more philosophical bent. While we do not claim to provide a full answer to the question, we offer an tentative interpretation of the notion of confidence in preference developed above, and some considerations on its relationship to other notions in the literature.

Let us begin by recalling three major positions in the debate over incomparability of value relations (see Chang (1997, 2002) for a presentation and references), which can be seen as loosely analogous to three sorts of positions in the literature on vagueness (see Keefe and Smith (1996) for a presentation and references). Basically, incomparability, like vagueness, could be either in the mind, in the world, or in language.

According to the "epistemic" position, there is no value incomparability, just ignorance. This is analogous to the position which claims that there is nothing more to vagueness than ignorance. So, for example, the betterness relation is perfectly determinate – there is a fact of the matter for any pair of objects that can fall under it which is better or whether they are equal – but we do not (and perhaps never will) know some such facts. Similarly, under such an approach to vagueness, there is, for each person, a fact of the matter whether he (or she) is bald, we just might not (and perhaps never will) know it. Let us call this epistemic position, insofar as it applies to value relations, *imprecision*.

According to the "ontic" position, value relations may be incomplete: it is a fact that, for some pairs of alternatives, say x and y, it is false both that x is (weakly) better than y, and that y is (weakly) better than x. This can be seen as analogous to the position that there are vague objects, that is, objects that have indeterminate properties; mountains or clouds

20. Note however that there would be no axiomatic difference, in terms of the Theorem VI.1, between the two cases.

are sometimes taken as examples (they purportedly have indeterminate boundaries). In these cases, no amount of information, or indeed reformulation in another language, would be able to resolve the indeterminacy, for it is not due to ignorance or imperfections in the language; it is a "hard" fact. Let us call this "ontic" position, insofar as it applies to value relations, (ontic) *indeterminacy*.

According to the "semantic" position, value relations may be semantically indeterminate, but not necessarily incomplete. This can be seen as analogous to doubtless the most popular position regarding vagueness, namely that it is a property of language. So, difficulty in attributing the predicate 'bald' to a particular person is not due to our ignorance, or the fact that this person is some sort of vague object, but rather down to the ways that the term 'bald' in our language is used and the way it relates to (potentially precise) objects in the world. Likewise, this position admits that whereas it may occur that the statements 'x is better than y' and 'y is better than x' are both not true, this does not imply that both of these statements are false or that the betterness relation is indeterminate in the sense specified above. Given its predominance in the literature on vagueness, let us call this position, insofar as it applies to value relations, (semantic) *vagueness*.

One way of transposing these positions onto the question of completeness of preferences is to consider the relation 'the agent prefers ... to ...'. It will be important that this notion of preference (as used in our language, or more specifically in the language of an economist or behavioural scientist) concerns attributions of preference to someone other than the person who is attributing. Comparison with the aforementioned literatures inspires a rudimentary taxonomy of theories in economics that drop or weaken the standard completeness assumption. The preference for flexibility models, where the current agent may be uncertain about his future preferences (see above and Section VI.3.2) correspond to imprecision. Theories of incomplete preference (see above and Section VI.1.1) would naturally seem to correspond to indeterminacy. Finally, given the relationship to the fuzzy theory of vagueness, fuzzy theories of preference might be most naturally construed in terms of semantic vagueness. Of course this is very crude, and several theories might be understood as straddling the boundaries between these categories.[21]

21. For example, whilst the interpretation of the fuzzy theory as vagueness is suggested by the voluntary reference to the literature on vagueness in, say, Salles (1998) and Piggins and Salles (2007), there are occasional hints that an epistemic reading is also intended (see for example what Sen calls the "pragmatic reason for incompleteness" in the passage cited by Salles (1998)).

That said, one suspects that much of the intuition for the latter two families of theories is not drawn from our use of the term ‘preference’ when describing others, but rather from the feeling *we* have that *our* preferences are undetermined or fuzzy. Although under the strictures of the modern paradigm in economics such considerations are not pertinent for comparisons among theories, they do often play an important motivating role in practice. Hence, in order to gain an insight into the conception behind the current proposal, and the comparison with others, it is perhaps worth considering the subjective, first-person point of view. As concerns the difficulties for the completeness property which are involved in the consideration of one’s own, current preferences, there is good reason to think that they are neither imprecision, nor indeterminacy, nor vagueness (in the senses described above), nor any combination of the three, but something entirely different.

In a word, the contrast between imprecision on the one hand, and indeterminacy and vagueness on the other rests on a sharp distinction between an epistemic agent and a set of facts that are independent from that agent. But it is far from clear that there is a neat separation between the relevant “facts” – the subject’s preferences at the moment of decision – and his beliefs about these “facts” – his beliefs about his preferences. Since, especially in deliberate decision making, one’s decisions are made on the basis of the preferences that one recognises in one way or another, what would it mean for one to have mistaken beliefs about one’s current preferences, that is, those which inform the actual choice being made? What would it mean to say that one’s actual preferences, at the moment in question, were not as experienced: that, for example, there were facts about them that were beyond one’s knowledge? If an agent couldn’t determine which he prefers out of x and y, what sense can be made of the question of whether he *really* does prefer x to y but doesn’t know it, or whether he *really* has indeterminacy of preference between x to y?

Given the connection between preference and choice, these worries are naturally related to the debate on the possibility of having beliefs about what one will choose, during the process of deliberation of that very choice (see for example Levi (1997); Joyce (2002); Rabinowicz (2002)). A moral that can be drawn from this debate is that, at moments of deliberation, beliefs about the outcomes of the deliberation lack many of the properties one usually associates with beliefs; so much so that one might wonder to what extent one can talk of belief at all. The suggestion is that, at the moment of choice, the notion of belief about preferences collapses in a similar way. But without it, one cannot say what it is for the purported object – preferences – to be inherently vague or indeterminate, rather than determinate but the subject of imprecise or uncertain beliefs.

If this is correct, then the boundaries between imprecision on the one hand, and indeterminacy and vagueness on the other, insofar as they apply to one's own current preferences, collapse. Whilst this is sufficient for the conclusions drawn below, let us note that the distinction between (ontic) indeterminacy of one's own preferences and (semantic) vagueness regarding them is even harder to defend. What would it mean for the preference relation used by the agent in his "private" language to be vague, rather than the "preferences themselves" being simply indeterminate?

If one is sensitive to these points, then it is evident that the notions of imprecision, indeterminacy and vagueness identified above are inappropriate to describe what may be lacking from our preferences at the moment of decision. A notion is required which does not separate the subject's attitude from what it is that is being assessed (and, moreover, from the conceptual framework in which it is being assessed). We suggest the term *confidence* to denote this notion. As such, confidence in preferences is understood as an intrinsic property of one's own, current preferences, in much the same way as vagueness may be thought of as an intrinsic property of the predicate 'bald' or incompleteness an intrinsic property of Schubert's unfinished symphony, and unlike, say, the property of being believed to be higher than 300m, which is not an intrinsic property of the Eiffel Tower. Nevertheless, the notion of confidence in preferences retains a doxastic aspect, without reducing to a fully fledged belief, because there is no solid distinction between a fully fledged belief and an independent object of belief in the case of one's own, current preferences.

The conceptual relationships with other accounts follow as a corollary. Assimilating confidence in preference to uncertainty about preference or imprecision in preference, or semantic vagueness or fuzziness in preference, or ontic indeterminacy in preference is not simply incorrect; it would be a category mistake. It makes no sense to speak of one's uncertainty about one's own current preferences in the same way as one speaks of one's uncertainty about the closing price of a particular company's shares, or about whether one thing is objectively better than another. Similarly, it makes no sense to speak of one's own current preferences being vague or fuzzy in the same sense as, say, the predicates 'bald' or 'poor' are.

It follows that it would be a mistake to assimilate the notion of confidence in preferences to any of the three positions in the philosophical literature on value relations described above,[22] or indeed to any application of the models discussed above in the study of value relations such as (objective) betterness or inequality. As noted, these positions depend on an assumed distinction between an epistemic agent considering the value relation and agent-independent facts about this relation. Of course, the cited literature tends to assume that there are such facts; the notion of confidence, as defined above, may have something to contribute, but only if one were willing to weaken this assumption.

From a modelling perspective, it follows that models drawn from the literature on uncertainty or vagueness must be justified from scratch if they are to be applied to (one's own) preferences, for this case is different from others in which they have previously been used. For example, those proposing models involving probabilities on the set of weak orderings cannot motivate them by interpreting them in terms of the agent's uncertainty about his current preferences and relying on the standard arguments for probabilities as representations of belief, for it is not (normal) belief which is at issue. Likewise, those proposing fuzzy models cannot motivate them by the intuition that our preferences are vague and relying on standard arguments for modelling vagueness in terms of degrees of truth, for it is not degrees of truth which are at issue. Of course, the recognition of the necessity for such justifications does not imply that they are not forthcoming.[23]

Let us conclude by emphasising that the last part of the discussion above is simply an attempt to situate the notion of confidence in preferences conceptually with respect to other notions of imprecision, indeterminacy, or vagueness. Economically speaking, the important part of the proposal is its operationalization in a precise formal model and fully axiomatised choice rule, carried out in Sections VI.1 and VI.2; comparison with other proposals can rely entirely on these aspects, without needing to enter into the conceptual intricacies.

22. Naturally, a similar point holds for other positions in that literature, such as the "parity thesis" defended by Chang (2002).

23. Note that some models may be motivated by considerations that are independent of the agent's first-person perspective on his current preferences, in which case the point made here does not apply. Examples include the use of probabilities in the preference for flexibility literature, which may be justified by standard arguments for probabilistic belief, because the beliefs in question do not concern the agent's current preferences but his future ones. Another example is the interpretation of probabilities in the random utility model as representing a random process, of which the agent may be unaware, that generates the utility function which he adopts and uses in his choice. When models are motivated but such considerations, they are representing a different phenomenon from the one that the model proposed here is concerned with.

VI.4 Social choice and confidence

With an eye to illustrating the interest of the notion of confidence in preferences proposed here, let us briefly consider an application to social choice. This discussion is not intended to be a comprehensive discussion of the potential importance of the notion of confidence for social choice, but rather a preliminary exploration of some possibilities.

The basic idea is that, if agents differ not only in their preferences but in their confidence in their preferences, then the latter factor and not solely the former can and often should be taken into account in the determination of the society's preferences. This makes sense: an agent's confidence in a preference (for x over y, for example) reflects how "sure" he is that he is "right" (by his own lights). Hence in aggregating the agents' preferences, it is not unreasonable to give those preferences in which an agent is more confident more bearing than those in which he is less confident. There are of course several ways in which this can be done; here we will consider only one.

As regards the setup, the set of alternatives, weak orderings and so on are as specified in Section VI.1.1. The set of members of the society (or voters) are assumed to be numbered, so the set of voters (which is not necessarily fixed) is some $V \subseteq \mathbb{N}$. Voters give not just their preferences but also their confidence in their preferences, which, as argued above, can be represented by an implausibility order. So a *profile* is a function $w : V \to \mathcal{I}$. The task is to determine a social preference ordering on the alternatives on the basis of each possible profile. Given that the agents' preferences are not necessarily fully determinate, we allow that the social preference ordering may be indeterminate; as noted in Section VI.1.1, this can be captured by representing it either by a quasi-ordering or by a set of weak orderings; here we use the latter option. The objects of study are thus functions which associate to each profile a subset of $\mathcal{P}$. We shall call such functions *confidence-adjusted social choice functions* (CASC), and denote them using the generic term f.

Following on from the intuition stated above, a natural CASC would be one which selects social preferences that, on aggregate, the members of the society are most confident in. Under one way of spelling out this idea, the CASC would aim to minimise the "total implausibility" of the social preferences (by the lights of the members of the society); this is like maximising the "total confidence" of society in the social value judgements. This is by no means the only way to go; we shall briefly discuss other options below.

Every weak ordering in $\mathcal{P}$ has a place in the implausibility order of each of the members of the society; this place can be "counted" by associating to the weak ordering its "rank" on the implausibility order $\leqslant$. Formally, the "rank" of a weak ordering R under an implausibility order $\leqslant$, $n_{\leqslant}(R)$, is defined as follows: $n_{\leqslant}(R) = \sup_{R'<R}(n_{\leqslant}(R')) + 1$, where the maximum over an empty set is taken to be -1. So the orderings at the bottom of the implausibility order (the "most plausible" ones) are of rank 0, those one rung up are of rank 1 and so on. The rank of a weak ordering can be thought of as a measure of the "distance" the ordering is from plausibility, according to the implausibility order in question. (Figure VI.2 in Section VI.1.1 makes this metaphor more vivid.)

A simple CASC which translates the idea that the social preference should be that which minimises its total implausibility is the "additive rank-based" CASC.

Definition VI.4. The *additive rank-based CASC* f is defined as follows: for any profile w, for any $R \in \mathcal{P}$,

(VI.1)
$$R \in f(w) \text{ iff } \sum_{v \in V} n_{w(v)}(R) \leqslant \sum_{v \in V} n_{w(v)}(R') \text{ for all } R' \in \mathcal{P}$$

The additive rank-based confidence-adjusted social choice function picks out the set of weak orderings whose total "implausibility", as summed over all the voters, is minimal (not greater than the total implausibility of any other orderings). In this sense, it could be thought of as maximising the total confidence in the social value judgements. Of course, although only a function yielding a set of weak orderings on alternatives has been defined, this can be easily extended to a definition of a function yielding an "social" implausibility order (that is, an order on the set of weak orderings).

Observe that the additive rank-based confidence-adjusted social choice function is none other than the Borda rule, applied to orderings over alternatives rather than to alternatives themselves. Thanks to this, we immediately have, borrowing a result from Young (1974), the following characterisation of this rule.[24]

Following Young, we say that a CASC f is *neutral* if, for σ a permutation of the set of orderings $\mathcal{P}$, and $\hat{\sigma}$ the induced permutation of profiles, $f(\hat{\sigma}(w)) = \sigma(f(w))$ for all profiles w. It is *consistent* if, for any w, w' profiles for disjoint voter sets V and V', then

24. Note that, though his theorem is stated for linear orderings, Young notes in the conclusion that it applies to weak orderings as well. Naturally, the case of weak orderings is that which is relevant here.

$f(w) \cap f(w') \neq \varnothing$ implies that $f(w) \cap f(w') = f(w + w')$ (where $w + w'$ is the profile on $V \cup V'$ which agrees with w on V and w' on V'). It is *faithful* if, for w a profile for one voter, $f(w)$ contains only the center of the voter's implausibility measure. Finally, a CASC f has the *cancellation property* if, whenever w is a profile such that, for any $R, R' \in \mathcal{P}$, the number of voters with $R < R'$ equals the number of voters with $R' < R$, then $f(w) = \mathcal{P}$. The following holds.

Theorem VI.4. *A CASC f is neutral, consistent, faithful and has the cancellation property if and only if it is the additive rank-based rule.*

As indicated, the conditions involved here are versions of standard conditions in the literature, and the reader is referred to the relevant papers (especially Young (1974)) for further discussion.

The purpose of these considerations is to give a flavour of possible applications of the notion of confidence to social choice. There are several directions that one could develop; let us just mention two.

First of all, the additive rank-based social choice rule is far from the only one, and others can be found and axiomatised in a way similar to that proposed, by exploiting the relation to voting theory. In fact, both the "additive" and the "rank-based" parts could be altered. For example, it is likely that an axiomatisation for a "maxmin rank-based" confidence-adjusted social choice function – which yields the set of those preference orders whose worst confidence ranking across voters is highest – can be obtained by using recent results on maxmin rules in voting theory (for example, Congar and Merlin (2011)). Or, to take another example, one might be able to develop and axiomatise an "additive importance-based" confidence-adjusted social choice function – where the total is taken not of the ranks of the weak orderings under the implausibility, but of the least importance levels which are associated to sets containing the weak orderings. Each suggestion appears to bring with it different issues, which may or may not be new. For example, the discussion of the relationship between the CASC proposed above and a minmax version may well mimic several classic debates in social theory, in particular the debate between utilitarianism and egalitarianism. By contrast, the comparison of rank-based and importance-based rules may well turn on the question of whether the agents' tolerances of choice in the absence of confidence (cautiousness coefficients) should be taken into account in the social preferences (as would be the case under the importance-based rule) or not.

Secondly, the sort of aggregation discussed above is "ordering-wise": it works with the order on the set of weak orderings $\mathcal{P}$. A further direction to explore is "judgement-wise" aggregation. As hinted in Section VI.1, an implausibility order represents whether the agent is more, less or equally confident in one value assessment (that alternative x is better than y by his lights) than in another (that alternative x' is better than y'). Under "judgement-wise" aggregation, one would not aggregate the rankings of the weak orderings under the implausibility order, but, say, the rankings of the value assessments on the order on value assessments generated by the implausibility order. This sort of aggregation may be interesting because the axioms would be expressed solely in terms of confidence in value assessments, and not in terms of orders on sets of weak orderings. Of course, there is a large literature on judgement aggregation which is relevant here (in particular, Dietrich and List (2010)). A particularly interesting question is the relation between judgement-wise and ordering-wise choice rules: is it the case, for example, that the set of value assessments endorsed by the result of an additive rank-based confidence-adjusted choice rule are those in which the total confidence is highest, as calculated by looking at the rankings of the assessments? This is, to our knowledge, an open question.

VI.5 Conclusion

People sometimes do not have preferences which are as determinate as the standard model would have us believe. Often, this is because people are not confident enough in some of the preferences they can be said to have. Of course, this may have implications for choice: people should not choose on the basis of preferences in which they are not sufficiently confident, if they can possibly avoid it.

This paper has made a start at bringing confidence in preferences into the field of choice theory. First of all, a representation of an agent's confidence in his preferences was developed, a notion of rationalisability of choice in terms of confidence in preference was proposed, and an axiomatisation of this notion was offered. The notion of rationalisability involves two main concepts which, to the knowledge of the author, have received relatively little attention in choice theory to date. Firstly, there is the concept of the importance of a choice, with the accompanying idea that the more important the choice, the more confident one needs to be in a preference to use it in one's choice. Secondly, there is the question of whether the agent can refuse to take a decision, or opt to defer, with the idea that this is the appropriate course of action when the choice is too important for the confidence he has in the relevant preferences. These notions, and their applications here, were discussed in detail. The technical and conceptual relationships with the existing economic and philosophical literature on indeterminacy were also examined.

Finally, in an attempt to indicate the relevance of the notion of confidence, a possible application to social choice was considered. A simple confidence-adjusted social choice function was proposed, based on the idea that the social preferences should be those in which the members of the society are, on aggregate, most confident. A simple axiomatisation was proposed for this rule, and some directions for future research were discussed.

Confidence in preferences has been given short shrift in choice theory to date. The author is confident that this should change.

VI.A Appendix

Proof of Theorem VI.2. Define the set of orderings $\mathcal{R}$ as follows: $R_i \in \mathcal{R}$ iff, for all $x, y \in X$, if $x \in c(\{x, y\})$, then xR_iy. First note that this set is well-defined. In particular $sing$ implies reflexivity and π implies the necessary transitivity: if $x \in c(\{x, y\})$ and $y \in c(\{y, z\})$, then by π and α, $x \in c(\{x, z\})$. Note also that this set is full: if R' agrees with the R in $\mathcal{R}$ wherever they all agree, then $R' \in \mathcal{R}$. Moreover, it is the unique full set.

It needs to be shown that this set of orderings generates c; consider $x \in S \subseteq X$.

Suppose $x \in c(S)$. Then, by α, $x \in c(\{x, y\})$ for all $y \in S$. So, xR_iy for all $y \in S$ and $R_i \in \mathcal{R}$, as required.

Suppose now that xR_iy for all $y \in S$ and $R_i \in \mathcal{R}$. Take an arbitrary enumeration of the elements of $S\backslash\{x\}$. We argue by induction that $x \in c(\{x, y_1, \ldots, y_n\})$ for all n. By hypothesis and definition of $\mathcal{R}$, $x \in c(\{x, y_1\})$. Suppose that $x \in c(\{x, y_1, \ldots, y_{n-1}\})$; by hypothesis and definition of $\mathcal{R}$, $x \in c(x, y_n)$; so by π, with $x = y$, $S = \{x, y_1, \ldots, y_{n-1}\}$ and $T = \{x, y_n\}$, $x \in c(\{x, y_1, \ldots, y_n\})$. Hence $x \in c(S)$, as required.

If c never takes as value the empty set, for all $x, y \in X$, either x or y (or both) belong to $c(\{x, y\})$. There is thus only one relation R such that for all $x, y \in X$, xRy iff $x \in c(\{x, y\})$: so the $\mathcal{R}$ constructed above is a singleton.

□

To prove Theorem VI.1, we first require the following Lemma.

Lemma VI.A.1. *Let choice* functions c_1 and c_2 be rationalised by full sets of orderings $\mathcal{R}_1$ and $\mathcal{R}_2$ respectively. If $x \in c_1(S)$ implies that $x \in c_2(S)$ for every $x \in S$ and every $S \subseteq X$, then $\mathcal{R}_1 \supseteq \mathcal{R}_2$.*

Proof. By construction of the rationalising sets of orderings in the proof of Theorem VI.2. The construction implies that $R \in \mathcal{R}_i$ if and only if, for all $x, y \in X$, if $x \in c_i(\{x, y\})$ then xRy (for $i = \{1, 2\}$). However, for every $R \in \mathcal{R}_2$, we have that, for all $x, y \in X$, if $x \in c_1(\{x, y\})$ then, by hypothesis, $x \in c_2(\{x, y\})$, and so xRy; it follows that $R \in \mathcal{R}_1$, as required. □

Proof of Theorem VI.1. The "only if" direction is straightforward to check. We consider here the "if" direction.

For any $i \in I$, note that $c(\bullet, i)$ is a function from sets of alternatives to sets of alternatives; it is a choice* function because the image may be empty. We will denote this function by c_i in what follows.

α^*, π^* and $sing^*$ imply that, for every $i \in I$, c_i satisfies α, π and $sing$. Theorem VI.2 implies that for each $i \in I$, c_i is rationalisable by a unique full set of weak orderings $\mathcal{R}_i$. Moreover, by Consistency and Lemma VI.A.1 (below), if $i \preccurlyeq i'$, then $\mathcal{R}_{i'} \subseteq \mathcal{R}_i$. Define $\leqslant$ as follows: $R \leqslant R'$ iff for all i such that $R' \in \mathcal{R}_i$, $R \in \mathcal{R}_i$. It is straightforward to check that this is complete, transitive and reflexive; ie. that it is an implausibility order. Define D by: $D(i) = \mathcal{R}_i$.

The representation of c by $\leqslant$ and D follows immediately from the construction. Also, by construction, $\leqslant$ is full, and any coarser full relation would fail to rationalise c; the uniqueness of D follows by construction. Consider finally the clause regarding centering. Centering implies that for every $S \subseteq X$, there exists $i \in I$ such that $c_i(S)$ is non-empty; by Consistency, there exists $i^* \in I$ such that, for all $S \subseteq X$, $c_{i^*}(S)$ is non-empty. By the final clause in Theorem VI.2, c_{i^*} is a singleton. This is the center of $\leqslant$.

□

Proof of Theorem VI.3. Define the implausibility order as in the proof of Theorem VI.1, using the part of c defined on menus containing †. It follows from the reasoning in the proof of that theorem that, for all $x \in X$ and all $S \subseteq X$, $x \in c(S \cup \{\dagger\}, i)$ iff $x \in \sup(S, D(i))$. By Deferral, if $\sup(S, D(i))$ is non-empty then $c(S \cup \{\dagger\}, i) = \sup(S, D(i))$; by the fact that the choice function always yields non-empty sets, it follows that if $\sup(S, D(i))$ is empty then $c(S \cup \{\dagger\}, i) = \{\dagger\}$. Moreover, if $\sup(S, D(i))$ is non-empty, then by $\alpha^\dagger$, $c(S, i) = c(S \cup \{\dagger\}, i) = \sup(S, D(i))$. Finally, it can be seen that if $\sup(S, D(i))$ is empty, then $x \in c(S, i)$ if and only if $x \in \sup(S, D(j))$ for all j such that $\sup(S, D(j))$ is non-empty. For if not, then there is a $j \in I$ such that $\sup(S, D(j))$ is non-empty but does not contain $x \in c(S, i)$. So $x \notin c(S, j)$ but $y \in c(S, j)$ for some y. Since $\sup(S, D(j))$ is non-empty and $\sup(S, D(i))$ is empty, $c(S \cup \{\dagger\}, j)$ is not contained in $c(S \cup \{\dagger\}, i)$, so, by Consistency†, it is not the case that $i \preccurlyeq j$. But, given $j \preccurlyeq i$ and $x \in c(S, i)$, Consistency† implies that $x \in c(S, j)$ contrary to the assumption.

Uniqueness follows from construction, as in the proof of Theorem VI.1.

□

Bibliography

Aumann, R. J. (1962). Utility Theory without the Completeness Axiom. *Econometrica*, 30(3):445–462.

Basu, K. (1984). Fuzzy revealed preference theory. *Journal of Economic Theory*, 32(2):212–227.

Broome, J. (1991). *Weighing Goods*. Basil Blackwell, Oxford.

Broome, J. (1997). Is Incommensurability Vagueness? In Chang, R., editor, *Incommensurability, Incomparability and Practical Reason*. Harvard University Press.

Chang, R. (1997). Introduction. In Chang, R., editor, *Incommensurability, Incomparability and Practical Reason*. Harvard University Press.

Chang, R. (2002). The Possibility of Parity. *Ethics*, 112(4):659–688.

Congar, R. and Merlin, V. (2011). A Characterization of the Maximin Rule in the Context of Voting. *Theory and Decision*, 72(1):131–147.

Danan, E. (2003). A behavioral model of individual welfare. Technical report, Université Paris 1.

Dietrich, F. and List, C. (2010). The aggregation of propositional attitudes: towards a general theory. *Oxford Studies in Epistemology*, 3.

Donaldson, D. and Weymark, J. A. (1998). A Quasiordering Is the Intersection of Orderings,. *Journal of Economic Theory*, 78(2):382–387.

Dutta, B., Panda, S. C., and Pattanaik, P. K. (1986). Exact choice and fuzzy preferences. *Mathematical Social Sciences*, 11(1):53–68.

Eliaz, K. and Ok, E. A. (2006). Indifference or indecisiveness? Choice-theoretic foundations of incomplete preferences. *Games Econ. Behav.*, 56:61–86.

Fine, K. (1975). Vagueness, Truth and Logic. *Synthese*, 30:265–300.

Fishburn, P. C. (1998). Stochastic Utility. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, volume 1. Kluwer, Dordrecht.

Hill, B. (2013). Confidence and decision. *Games and Economic Behavior*, 82:675–692.

Joyce, J. (2002). Levi on Causal Decision Theory and the Possibility of Predicting one's on Actions. *Philosophical Studies*, 110:69–102.

Keefe, R. and Smith, P. (1996). Introduction. In Keefe, R. and Smith, P., editors, *Vagueness: A Reader*. MIT Press, Cambridge, MA.

Kreps, D. M. (1979). A Representation Theorem for 'Preference for Flexibility'. *Econometrica*, 47:565–576.

Levi, I. (1986). *Hard Choices. Decision making under unresolved conflict*. Cambridge University Press, Cambridge.

Levi, I. (1997). Rationality, Prediction and Autonomous Choice. In *The Covenant of Reason. Rationality and the Commitments of Thoughts*. CUP, Cambridge.

Luce, R. D. and Suppes, P. (1965). Preference, Utility and Subjective Probability. In Luce, R. D., Bush, R. R., and Galanter, E. H., editors, *Handbook of Mathematical Psychology*, volume 1. Wiley.

Mandler, M. (2009). Indifference and incompleteness distinguished by rational trade. *Games Econ. Behav.*, 67(1):300–314.

Morton, A. (1991). *Disasters and Dilemmas*. Blackwell, Oxford.

Moulin, H. (1985). Choice Functions Over a Finite Set: A Summary. *Social Choice Welfare*, 2:147–160.

Piggins, A. and Salles, M. (2007). Instances of Indeterminacy. *Analyse & Kritik*, 29(2):311–328.

Rabinowicz, W. (2002). Does Practical Deliberation crowd out Self-Prediction? *Erkenntnis*, 57:91–122.

Rabinowicz, W. (2008). Value Relations. *Theoria*, 74(1):18–49.

Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge, MA.

Salles, M. (1998). Fuzzy Utility. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory*, volume 1. Kluwer, Dordrecht.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Sen, A. K. (1970). *Collective Choice and Social Welfare*. Holden Day, San Francisco.

Sen, A. K. (1973). *On Economic Inequality*. Oxford University Press, Oxford.

Sen, A. K. (1993). Internal Consistency of Choice. *Econometrica*, 61(3):495–521.

Sen, A. K. (1997). Maximization and the Act of Choice. *Econometrica*, 65(4):745–779.

Williamson, T. (1994). *Vagueness*. Routledge, London.

Young, H. P. (1974). An axiomatization of Borda's rule. *J. Econ. Theory*, 9(1):43–52.

VII Incomplete preferences and confidence

Abstract

A theory of incomplete preferences under uncertainty is proposed, according to which a decision maker's preferences are indeterminate if and only if her confidence in the relevant beliefs does not match up to the stakes involved in the decision. We use the model of confidence in beliefs introduced in Hill (2013), and axiomatise a class of models, differing from each other in the appropriate notion of stakes. Comparative statics analysis can distinguish the decision maker's confidence from her attitude to choosing in the absence of confidence. The model naturally suggests two possible strategies for completing preferences, and hence for choosing in the presence of incompleteness. One strategy respects confidence – it relies only on beliefs in which the decision maker has sufficient confidence given the stakes – whereas the other goes on hunches – it relies on all beliefs, no matter how little confidence the decision maker has in them. Axiomatic characterizations are given for each of the strategies. Finally, we consider the consequences of the model in markets, where indeterminacy of preference translates into refusal to trade. The incorporation of confidence adds an extra friction, beyond the standard implications of non-expected utility models for Pareto optima. [1]

Keywords: Incomplete preferences; confidence; multiple priors; choice under incomplete preferences; absence of trade.

JEL classification: D81, D01, D53.

1. This is a slightly modified version of the HEC Paris Research Paper No. ECO/SCD-2014-1051.

VII.1 Introduction

VII.1.1 Motivation

Incomplete preferences have been increasingly recognized as of importance. Appeals to the weakening of the completeness axiom – which demands that for every pair of options, the decision maker has a weak preference for one over the other – have been made both in the name of 'psychological realism' (Aumann, 1962; Dubra et al., 2004; Danan, 2003b; Galaabaatar and Karni, 2013) and on the basis of normative considerations (Aumann, 1962; Bewley, 2002). Moreover, incomplete preferences have proved invaluable in the development of alternative models of choice, such as those incorporating a tendency to stick to the status quo (Bewley, 2002; Masatlioglu and Ok, 2005). Incomplete preferences naturally arise is multi-agent settings, where the preferences of a group, or those drawn from group members' beliefs or utilities may naturally be incomplete (Dubra et al., 2004). As a final example, objectively rational preferences in the sense of Gilboa et al. (2010) – those preferences for which the decision maker can convince others of their correctness, by a form of proof for example – are naturally incomplete.

Standard approaches to modeling incomplete preferences often proceed, roughly speaking, by dropping the completeness axiom whilst retaining the other standard axioms, and replacing a single function or measure in the model by a set. For instance, in decision under uncertainty, the benchmark unanimity multi-prior model proposed by Bewley (2002) retains all standard Anscombe-Aumann (1963) axioms for subjective expected utility except completeness, and replaces the single probability measure in the representation by a set of probability measures. In particular, it retains the independence axiom.

However, under all of the interpretations mentioned above, there appear to be cases where the standard independence axiom is violated. Consider a decision maker who is faced with choices between bets on the color of the next ball drawn from an urn containing only black and white balls, as shown in Figure VII.1. For simplicity, suppose that the bets are given in dollars and the decision maker has linear utility.[2] She is told neither the proportion nor the number of balls in the urn, but she has observed fifteen draws (with replacement), nine of which were black and the rest of which were white. It does not seem

2. Alternatively, one could read the bets as given in utils, and as corresponding to the appropriate mixtures of corresponding dollar bets in the standard way; eg. f is the mixture $\frac{1}{100000}g + \frac{99999}{100000}0$.

Figure VII.1 – Bets ('M' stands for 'million')

	Colour of ball drawn from urn	
	Black	White
f	15	-10
0	0	0
g	1.5 M	-1 M
f^n	15 $\times n$	-10 $\times n$

implausible that there are decision makers who prefer f to 0 given this information, whilst being indeterminate in their preference between g and 0. Certainly, from a normative point of view, it is not unreasonable to hold a preference between the first pair of bets while not having a determinate preference between the second pair, given the weakness of the information and the stakes involved. Even from the point of view of objective rationality, there is a 'statistical argument' for preferring f over 0 – based, for example, on a classical hypothesis test with a weak significance level (eg. 10%)[3] – whereas there is no objectively rational preference between g and 0 – in the situation where more is at stake, arguably more stringent standards of proof, such as tougher significance levels, are required, and the data do not support any conclusions at such levels. Analogous cases exist for the group interpretation of incomplete preferences: for example, if there is agreement between two leading urn-experts that the proportion of black balls is $\frac{1}{2}$, but a large disagreement in the community as a whole on the proportion of black balls, it is does not seem unreasonable for the group to form a preference between f and 0 without forming one between g and 0. Since independence implies that there is preference for f over 0 if and only if there is preference for g over 0, it is violated in these examples.

Reinterpreting the event that the ball is black to be the success of a new technology, for example, and the observations to be suggestive yet inconclusive findings, it is clear that there are real-life cases where this sort of preference pattern is exhibited. On the basis of limited grounds (be they scarce information, a weak argument or a few members of the

3. Explicitly, a one-sided classical statistical test rejects the hypothesis that the proportion of black balls is 0.4 at the 10% significance level, and for probabilities of black above 0.4, f has a higher expected utility than 0.

group), decision makers may be ready to form preferences when the decision is relatively unimportant, but cannot do so when there is more at stake. Our proposed diagnosis is that the standard models of incomplete preferences (in terms of sets of probability measures, for example) overlook the fact that decision makers can be more or less *sure* of their beliefs. The examples given above suggest that *how* sure the decision maker is in a belief may be related to her preferences over options for which this belief is relevant. These appear to be cases where determinate preferences are formed on the basis of beliefs in which the decision maker is not entirely sure in some situations – in particular, when little is at stake in the decision – whereas there are other situations – when the decision is more important, for example – in which she may need to be more sure of her beliefs to avoid indeterminacy.

Beyond the case of individual preferences, this diagnosis also applies to other interpretations of incompleteness. A group – a firm, for example – may be said to hold a belief if appropriate subgroups – the directorship, or larger or smaller collections of committees – are unanimous on the belief, and it is natural to say that the group is more sure of a belief shared by all stakeholders than one shared only by the board of directors. Moreover, whilst agreement among directors may be sufficient for some decisions, more group consensus would seem a reasonable prerequisite when the stakes are high. Similarly, there are more or less stringent standards of proof (corresponding, for example, to the strength of statistical tests, the degree of experimentation, the rigor of argument, and so on) that could be involved in objectively rational preferences. Any notion of objectively rational preference will have to specify the appropriate standards: a natural way of doing so is by demanding more rigorous proof when the choices are more important.

The aim of this paper is to propose a theory of decision under uncertainty that, whilst deviating as little as possible from the standard models of incomplete preference, incorporates the decision maker's confidence in her beliefs. The theory can be summed up in the following principle: one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form the preference does not match up to the stakes involved in the choice. Like the standard Bewley model, indeterminacy of preferences is driven solely by the decision maker's beliefs; it is tacitly assumed that the decision maker is fully confident in her utilities.

As concerns behavioral properties, note that in the context of incomplete preferences, independence applied to the preference $f \succ 0$ and the acts g and 0 (Figure VII.1) in fact implies two distinct things: on the one hand, there is a determinate preference between g and 0; on the other hand, this preference goes in the appropriate direction ($g > 0$). The examples above only conflict with the former condition, not the latter; however, it is the latter condition that is at the heart of the independence property. Hence it is natural to drop the former condition, retaining the latter: that is, to demand that the standard independence condition applies whenever the preferences involved are determinate. This is the appropriate weakening of independence for the model developed in this paper. Indeed, the other main axiomatic difference from the Bewley multi-prior model involves a similar weakening of transitivity: it applies whenever preferences are determinate, but indeterminacy is permitted in some cases where standard transitivity would have demanded determinate preference. [4] We take the mildness of these axioms to be an indication of the parsimony of this departure from the benchmark Bewley model of incomplete preferences under uncertainty.

The incorporation of confidence into an account of incomplete preferences has interesting consequences for the question of how to 'complete' preferences, a question that is pertinent under all the aforementioned interpretations, in particular when a decision must be taken. It allows the distinction between, and characterization of, two strategies for preference completion. One respects confidence, insofar as it only allows the decision maker to use beliefs in which she has sufficient confidence given the stakes involved in the decision. A government who bases its climate policy on 'full scientific certainties', however scarce they may be and ignoring the less well-established opinions of experts, adopts this strategy. The other strategy goes on hunches, insofar as it allows the decision maker to mobilize all her beliefs – even those in which she has little confidence – when she is forced to choose. An entrepreneur who undertakes a venture on the basis of her 'gut feeling' that it will work, without being strongly convinced of its success, is adopting this strategy. The distinction between these strategies, though pre-theoretically reasonable and potentially pertinent to the understanding of real-life decisions, has not yet been identified in the literature, to our knowledge.

4. The need for a weakening of transitivity can also be seen on the example above. It is not implausible, in the light of similar considerations to those behind the preference for f over 0, that the decision maker prefers f^{n+1} to f^n for all n between 0 and 99 999 (recall that she has linear utility). Transitivity would imply that she prefers g over 0, and hence is violated. See Section VII.3.1 for further discussion.

VII.1.2 Overview of proposal and results

To develop a theory of decision based on the maxim that one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form the preference does not match up to the stakes involved, we use several notions introduced in Hill (2013). The decision maker's confidence in her beliefs is represented by a *confidence ranking*: that is, a nested family of sets of probability measures. Each set in the family is interpreted as corresponding to a level of confidence. A probability judgement [5] that applies under every probability measure in the set is one that the decision maker holds to the corresponding level of confidence. Larger sets in the family correspond to higher levels of confidence; fewer probability judgements (or beliefs) are held to those levels of confidence. Whilst proposed as a representation of individual beliefs, confidence rankings have a natural interpretation for groups. Each probability measure can be thought of as the beliefs of a member of the group, and each level as corresponding to a level in the group's hierarchical structure (eg. in a country, one level will contain cabinet ministers, another will contain members of the government, another all elected representatives, and so on). The probability held at a particular level is one that is shared unanimously by all group members who have at least the rank corresponding to that level.

The decision maker's attitude to choosing in the absence of confidence is represented by a *cautiousness coefficient*: a function that associates to each decision a set of probability measures in the confidence ranking. It can be thought of as assigning the level of confidence required in beliefs for them to play a role in the decision in question. It is defined in such a way that it picks out the same set of probability measures for all decisions having the same stakes; as such, it can be thought of as picking out the appropriate level of confidence entirely on the basis of the stakes involved in the decision. The confidence ranking and cautiousness coefficient are subjective elements in the model; for extended discussion and defense of these notions and their interpretations, see Hill (2013) and Section VII.3.3 below.

We represent the notion of stakes by a *stakes relation* on the set of decisions, which orders decisions according to whether the stakes involved in them are higher or lower. In this paper, we take the view that the notion of stakes is a feature of the decision model, different notions of stakes yielding different decision models. One can imagine several dif-

5. By probability judgement, we mean a statement concerning probabilities, such as 'the probability of event A is greater than p'.

ferent notions of stakes (see Appendix VII.A), and we are non-committal about the proper stakes relation to use. Hence we exogenously assume a stakes relation satisfying some basic properties. The results hold for any such stakes relation, and are interpreted as providing *inter alia* an axiomatic treatment of a class of models, which differ only in the notion of stakes used. An alternative perspective would be to consider the notion of stakes to be subjective to the decision maker; Hill (2014) adopts such a perspective in a different but related framework, and can be thought of as providing a behavioral foundation for the stakes relation.

In the main result of the paper, we give necessary and sufficient conditions for a representation of the following form: for all Anscombe-Aumann acts f and g, $f \preceq g$ if and only if

(VII.1) $$\sum_{s\in S} u(f(s)).p(s) \leqslant \sum_{s\in S} u(g(s)).p(s) \quad \text{for all } p \in D((f,g))$$

where u is a von Neumann-Morgenstern utility function and D is a cautiousness coefficient for a confidence ranking Ξ. To understand this representation, note firstly that $D((f,g))$ is the set of probability measures associated with the choice between acts f and g, and depends entirely on the stakes involved in this choice. It indicates the confidence level associated to those stakes; the only probability judgements in which the decision maker is confident enough for them to play a role in the decision are those which hold for all probability measures in $D((f,g))$. Under representation (VII.1), g is weakly preferred to f if and only if, based only on these probability judgements, the decision maker can conclude that the expected utility of g is at least as high as that of f. So if, on the basis of these probability judgements, the decision maker can conclude neither that g has expected utility at least as high as f nor that f has expected utility at least as high as g, then she has no preference between them. In other words, her preferences over a pair of acts are indeterminate if she is not confident to the degree required by the stakes involved in the decision that one act is at least as good as the other.

This model permits comparative statics analysis of the decision maker's decisiveness – her propensity to exhibit determinate preference – and supports a behavioral definition and characterization of the decision maker's confidence in her preferences. This analysis vindicates the interpretation, mooted above, of the cautiousness coefficient as capturing her attitude to choosing in the absence of confidence.

Moreover, as intimated above, the model supports an analysis of two families of completion strategies for preferences, which differ in the beliefs that the decision maker allows herself to rely on when completing her preferences. The difference between the strategies is naturally captured by the level of the confidence ranking used in the formation of complete preferences. Behavioral conditions that yield simple representations for each of the strategies are proposed.

Finally, a standard interpretation of indeterminacy of preferences in market settings (dating back at least to Bewley (2002)) is in terms of reluctance to trade. What implications does the incorporation of confidence according to representation (VII.1) have in such settings? We show that it adds a friction absent under other non-expected utility or incomplete preference models of decision under uncertainty, with consequences for the difficulty of attaining a Pareto optimum via Pareto-improving trade.

The basic notions of the model are introduced and formally defined in Section VII.2. The representation result and comparative statics of decisiveness are given in Section VII.3. In Section VII.4, we consider the question of how to complete one's incomplete preferences. In Section VII.5, we consider the consequences of the model in markets under uncertainty. Proofs of all results and other material are to be found in the Appendices.

In the rest of the Introduction, we discuss the relation to existing literature.

VII.1.3 Related literature

Bewley (2002) was the first to axiomatise a 'unanimity' representation of an incomplete preference relation by a set of probability measures, according to which there is a preference between acts if the expected utilities of the acts lie in the appropriate relation for all the probability measures in the set. Technically, our representation is closer to the unanimity representation used by Gilboa et al. (2010), who take the weak rather than the strict preference relation as primitive.[6] The unanimity model cannot capture differing degrees of confidence, and hence it does not have the richness to capture the effect of the stakes involved in a choice on the degree of confidence required of beliefs to play a role in it, and hence on determinacy of preferences. Representation (VII.1) can thus be thought of as

6. The representation in Gilboa et al. (2010) differs from representation (VII.1) above by replacing $D((f, g))$ with a fixed set of probability measures; the representation in Bewley (2002) differs moreover in replacing the weak preferences and orders by strict preferences and orders.

a generalisation of the unanimity representation, replacing a single fixed set of probability measures by a family of sets, where the set of measures used varies depending on the stakes involved in the decision.

Representation (VII.1) belongs to a family of decision models that represent the decision maker's state of belief by a confidence ranking and are based on the idea that different sets of probability measures may be used in the evaluation of options, according to the stakes involved. This family was introduced and motivated in Hill (2013). There it was noted that members differ along two dimensions: the decision rule which determines preferences on the basis of a set of probability measures and a utility function, and the notion of stakes. In this perspective, the current paper can be thought of as complementary to Hill (2013), exploring different fragments of the family introduced there. Theorem VII.1 axiomatises a family of models that take the unanimity decision rule (as opposed to the maxmin expected utility rule, as in Hill (2013)) and any notion of stakes that generates a stakes relation satisfying some basic properties (as opposed to a particular notion of stakes, as in Hill (2013)). Moreover, Theorem VII.2 (ii) can be thought of as providing foundations for another class of models belonging to the same family; namely the class using the maxmin expected utility decision rule and any notion of stakes of the sort described above. [7]

Nau (1992) has proposed a theory of incomplete preferences which is similar to representation (VII.1) in content and motivation. Besides the difference in framework (he uses the de Finetti framework), presentation (he uses confidence-weighted upper and lower conditional probabilities on random variables) and conceptualisation (the distinction between stakes and confidence is not fully brought out, and the notion of cautiousness coefficient is absent), he assumes a particular notion of stakes, whereas we do not. As discussed in Appendix VII.A, there is a sense in which his theory is a member of the family axiomatised here.

Faro (2013) has proposed an extension of Bewley's representation by incorporating a real-valued function on the space of probabilities measures, in a way inspired by the variational preferences model of Maccheroni et al. (2006). Lehrer and Teper (2011) have proposed a representation involving sets of sets of probability measures, where an act is

7. Note that the model considered in Hill (2013), although it uses the same decision rule, does not belong this class. There the stakes are a function of a single act, whereas here the stakes are a function of the decision (that is, the pair of acts on offer); see Section VII.2.2.

preferred to another if it has a higher expected utility for all probability measures in at least one of the sets. There are significant differences from the current proposal in the representation of preferences (for example, the notion of the stakes plays no role in these models), the concepts involved and how they are modeled, and the behavioral properties. In particular, these models employ more severe weakenings of transitivity than used here, and indeed allow preference cycles whilst representation (VII.1) does not; moreover, Faro (2013) employs a more severe weakening of independence. Faro (2013) also considers the relation to the variational preferences and maxmin EU models, in a manner similar to our treatment of the question of preference completion in Section VII.4.

Minardi and Savochkin (2013) propose a representation of a graded preference relation in terms of a capacity over a set of probability measures, where the 'strength of' or 'confidence in' the preference for an act is equal to the measure of the set of probability measures for which the expected utility of the act is greater. Their graded preference relation is a binary relation over pairs of acts, and hence is reminiscent of the confidence-in-preferences relation introduced in our Section VII.3.3, with the notable difference that whilst Minardi and Savochkin (2013) assume this relation as a primitive, here it is defined from (ordinary) preferences over acts (Definition VII.2, Section VII.3.3). Analogous points to those made in the previous paragraph appear to apply to the comparison with this model: for example, notwithstanding the difference in framework, their transitivity condition (Weak Transitivity) seems to be closer to the weakening of transitivity used in Faro (2013); Lehrer and Teper (2011) than to the one used here.

Seidenfeld et al. (1995); Nau (2006); Ok et al. (2012); Galaabaatar and Karni (2013) have explored extensions of Bewley's representation involving sets of probabilities and sets of utilities. The behavioral points made above, in particular concerning the independence axiom in the presence of incompleteness, continue to hold for these models. They plead in favour of the incorporation of confidence in beliefs and confidence in utilities; this is left as a topic for future research. Hill (2009) proposes a model of confidence in preferences that retains the same basic intuition as the models of confidence of belief used here, and applies it in the context of choice under certainty.

The discussion of the completion of incomplete preferences is technically related to Gilboa et al. (2010), Kopylov (2009) and Nehring (2009), who provide representation results on pairs of binary relations, where one is complete, the other is represented according to the unanimity representation described above, and they are represented by related or identical sets of probability measures. Although Gilboa et al. (2010), like us, propose a representation of the second relation in terms of the maxmin expected utility representation (Gilboa and Schmeidler, 1989), Kopylov (2009) and Nehring (2009) consider special and more general cases respectively. All these authors work with single sets of probability measures, rather than confidence rankings, and hence cannot capture the distinction between the two strategies presented in Section VII.4. A final, axiomatic, difference is that they all impose axioms other than completeness on the complete preference relation (generally, at least transitivity and continuity), whereas we only impose completeness. Accordingly, the main new axiom used here (Stakes Benchmark on Certainty, Section VII.4) is stronger that the main ‘connecting’ axioms used, for example, by Gilboa et al. (2010) (Consistency and Caution or Default to certainty). In the case of a degenerate confidence ranking (containing a single set of probability measures), our Theorem VII.2 thus provides a new axiomatisation of the representation obtained in their Theorems 3 and 4.

Finally, the interpretation of indeterminacy of preference in terms of sticking to a status quo option used in Section VII.5 has been considered by Bewley (2002), under the name of the ‘inertia assumption’. Bewley (1989) was the first to consider consequences for trade, and Rigotti and Shannon (2005) undertake a thorough analysis of markets involving decision makers with unanimity preferences. Billot et al. (2000), Rigotti et al. (2008) and Ghirardato and Siniscalchi (2014) consider markets involving decision makers with complete non-expected utility preferences.

VII.2 General preliminaries

VII.2.1 Setup

Throughout the paper, we use the standard Anscombe-Aumann framework (Anscombe and Aumann, 1963), as adapted by Fishburn (1970). Let S be a non-empty finite set of states; subsets of S are called *events*. $\Delta(S)$ is the set of probability measures on S, endowed with the Euclidean topology. X is a nonempty set of outcomes; a *consequence* is a probability measure on X with finite support. $\Delta(X)$ is the set of consequences. *Acts* are functions from states to consequences; $\mathcal{A}$ is the set of acts. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in [0,1]$, the mixture $\alpha f + (1-\alpha)h$

is defined by $(\alpha f + (1-\alpha)h)(s,x) = \alpha f(s,x) + (1-\alpha)h(s,x)$. We write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$.

The decision maker's preferences over $\mathcal{A}$ are represented by a binary relation $\preceq$. $\sim$ and $\prec$ are the symmetric and asymmetric components of $\preceq$, and $\asymp$ is the 'determinate preference' relation, defined as follows: $f \asymp g$ iff $f \preceq g$ or $f \succeq g$. So $f \not\asymp g$ means that the decision maker does not have a determinate preference between f and g.

VII.2.2 Stakes

As mentioned in the Introduction, a central idea in this paper is the importance of the stakes involved in a choice between two options for the preferences over them. We represent the notion of stakes by a *stakes relation* that, for each pair of (binary choice) decisions the decision maker may be faced with, specifies which has higher stakes or whether the stakes involved in them are the same. That is, a stakes relation $\leqq$ is a relation on the set of pairs of acts ($\mathcal{A} \times \mathcal{A}$) that is interpreted as follows: $(f,g) \leqq (f',g')$ means that the stakes involved in the choice between f and g are (weakly) lower than the stakes involved in the choice between f' and g'. $\equiv$ and $<$ are the symmetric and asymmetric components of $\leqq$, defined in the standard way. We shall be interested in stakes relations satisfying the following basic properties.

(Weak Order) $\leqq$ is reflexive, transitive and complete.

(Symmetry) For all $f,g \in \mathcal{A}$, $(f,g) \equiv (g,f)$.

(Extensionality) For all $f,f',g,g' \in \mathcal{A}$, if $f(s) \sim f'(s)$ and $g(s) \sim g'(s)$ for all $s \in S$, then $(f,g) \equiv (f',g')$.

(Continuity) For all $f,f',g,g',h \in \mathcal{A}$, the sets $\{(\alpha,\beta) \in [0,1]^2 |\ (f_\alpha h, g_\beta h) \geqq (f',g')\}$ and $\{(\alpha,\beta) \in [0,1]^2 |\ (f_\alpha h, g_\beta h) \leqq (f',g')\}$ are closed in $[0,1]^2$.

(Richness) For all $f,f',g,g' \in \mathcal{A}$ such that $f(s) \not\asymp g(s)$ for some $s \in S$ and (f',g') is not $\leqq$-minimal, there exists $h,h' \in \mathcal{A}$ and $\alpha,\alpha' \in (0,1]$ such that $(f_\alpha h, g_\alpha h) \leqq (f',g') \leqq (f_{\alpha'} h', g_{\alpha'} h')$.

Weak order states that the binary choices the agent may be faced with can be weakly ordered according to the stakes involved in them. We take this to be a basic property of the notion of stakes, and accept it without discussion here.[8] Symmetry states that the stakes involved in a choice only depend on the alternatives available, irrespective of the order in which they are presented, and deserves no further discussion.

Extensionality says that all that counts for the stakes are the values of the consequences of the acts at the different states. If two acts are extensionally equivalent – that is, the decision maker is indifferent between the consequences of the acts at every state – then in virtually all formal theories of decision under uncertainty, they are treated (and evaluated) in exactly the same way. Extensionality says that whenever two pairs of acts are related in this way, they have the same stakes. Note that extensionality involves reference to the decision maker's preferences over constant acts. This is to be expected: the stakes involved in a choice depend on how 'good' and 'bad' the consequences of the various options are in particular instances, and of course, how 'good' and 'bad' they are depends on the decision maker's evaluation. Indeed, many plausible notions of stakes which come to mind make reference at least to the decision maker's utility function, if not also to part of her confidence ranking (see the examples in Appendix VII.A). It is thus entirely natural for the stakes relation to respect some aspects of her utility function.

On mixing a pair of acts with a third act, the stakes involved in a choice may change; Continuity says that this change is continuous in the degree of mixing. This seems reasonable: the stakes may be altered as one or both of the acts on offer are mixed with another act, but one would not expect the stakes to 'jump' as the mixture coefficient moves gradually from one value to another.

Richness is a technical property, which states that the stakes involved in a choice between different acts can be shifted as far up or down the stakes order as desired, by mixing the pair of acts appropriately. There is a sense in which (in particular in the presence of an independence axiom; see Section VII.3) the choice between f and g and the choice between $f_\alpha h$ and $g_\alpha h$ are the 'same' choice. Nevertheless, the stakes involved in these two choices may differ; to that extent, the latter choice can be thought of as a 'version' of the former choice, but at the stakes level corresponding to $(f_\alpha h, g_\alpha h)$ rather than (f, g). Hence, using such mixtures, one can consider the decision maker's preferences at different stakes levels. We will say that the decision maker prefers f to g at a certain stakes level if she

8. See Hill (2009, Section 3.1) for a discussion of the possibility of weakening the completeness assumption in a different but related framework.

prefers $f_\alpha h$ to $g_\alpha h$ for some α and h such that $(f_\alpha h, g_\alpha h)$ has that level of stakes. Richness simply states that for any non-trivial choice and non-minimal stakes level, there is a version of the choice, obtained by mixing with a third act, which has stakes above the level in question, and there is a version which has stakes below that level. The intuition is that mixing with a third act can change many of the properties of a pair of acts, and in particular the main properties that are relevant for the stakes involved in the choice between them. Under some notions of stakes, 'trivial' choices – between an act and itself, or, given extensionality, between an act and one which is extensionally equivalent to it – and the minimal stakes level have a special status; it could be, for example, that all and only trivial choices have minimal stakes (this is the case for some of the notions mentioned in Appendix VII.A). For this reason, richness does not apply to such choices and stakes levels.

As noted in Hill (2013, Section 4), there are many notions of stakes, each of which generates a stakes relation. Of course, the weaker the properties assumed for the stakes relation, the more general the results involving it. To illustrate the generality of the results in the following sections, we give some examples of notions of stakes satisfying the properties above in Appendix VII.A. We also discuss weakenings of the properties and extensions to allow dependence of the stakes on the status quo or the context that can be straightforwardly accounted for by appropriate modifications in the theorems below.

As mentioned in the Introduction, we adopt the perspective presented in Hill (2013), according to which the notion of stakes is an objective feature of the decision model: different notions of stakes yield different decision models. We thus assume throughout Sections VII.3 and VII.4 a stakes relation $\leqq$ satisfying the five properties above. The analysis and results hold for any stakes relation satisfying these properties; as such, they can be thought of as applying to a class of decision models, where the members of the class differ on the notions of stakes. Note that, to the extent that some axioms are formulated in terms of the stakes relation, different notions of stakes will correspond to different behavioral properties, and hence it is possible in principle to tell whether the decision maker is using a given stakes relation or not. Hill (2014), working with a different but related decision model, provides a representation in which the stakes relation (over general, rather than only two-element menus) is endogenous.

VII.2.3 Confidence ranking and cautiousness coefficient

We recall two notions that were introduced in Hill (2013). (The reader is referred to that paper for further discussion of these notions and their properties.)

Definition VII.1. A *confidence ranking* Ξ is a nested family of closed, convex subsets of $\Delta(S)$. A confidence ranking Ξ is *continuous* if, for every $C \in \Xi$, $C = \overline{\bigcup_{\Xi \ni C' \subsetneq C} C'} = \bigcap_{\Xi \ni C' \supsetneq C} C'$. It is *centered* if $\bigcap_{\mathcal{C} \in \Xi} \mathcal{C}$ is a singleton. [9]

As mentioned in the Introduction, confidence rankings represent decision makers' beliefs, and in particular their confidence in probability judgements. The sets in the confidence ranking can be thought of as corresponding to levels of confidence. The higher the level of confidence in question, the larger the set: this translates the fact that one is confident of fewer probability judgements to that level of confidence. The convexity and closedness of the sets of probability measures in the confidence ranking are standard assumptions for decision rules involving sets of probabilities. The continuity property guarantees a continuity in one's confidence in probability judgements: it ensures, for example, that one cannot be confident up to a certain level that the probability of an event A is in $[0.3, 0.7]$ and then only confident that the probability is in $[0.1, 0.9]$ at the 'next' confidence level up. All confidence rankings considered in this paper will be continuous, the property often being assumed in discussion without being explicitly mentioned. Finally, a decision maker with a centred confidence ranking is one who, if forced to give her best estimate for the probability of any event, could come up with a single value (and these values satisfy the laws of probability), although she may not be very confident in it. We do not assume confidence rankings to be centred in general; in the representation result, there is an axiom that is necessary and sufficient for the confidence ranking to be centred, so this property receives a behavioural characterisation.

The second notion required is that of a *cautiousness coefficient* for a confidence ranking Ξ, which is defined to be a surjective function $D : \mathcal{A} \times \mathcal{A} \to \Xi$ that preserves $\leqq$; that is, such that for all $(f, g), (f', g') \in \mathcal{A} \times \mathcal{A}$, if $(f, g) \leqq (f', g')$, then $D((f, g)) \subseteq D((f', g'))$. D assigns to any pair of acts the level of confidence that is required in probability judgements in order to be used in the choice between the acts. This level of confidence corresponds to the appropriate set of probability measures in the confidence ranking. Preservation of

9. For a set X, $\overline{X}$ is the closure of X. Note that the union of a nested family of convex sets is convex.

the stakes relation $\leqq$ implies that D assigns a confidence level to a choice solely on the basis of the stakes involved in that choice, and is faithful to the intuition that the higher the stakes, the higher the confidence required in probability judgements for them to play a role in the choice. Surjectivity of D basically attests to the behavioral nature of the confidence ranking: it implies that for each set of probability measures in the ranking, there will be a level of stakes, and hence a choice, for which it is the relevant set.

VII.3 A theory of incomplete preferences and confidence

In this section, we axiomatise and analyze a representation of incomplete preferences. Throughout this section and the next one, we assume a fixed stakes relation $\leqq$.

VII.3.1 Axioms

Consider the following axioms on $\preceq$.

Axiom VII-A1 (C-Completeness)**.** For all $c, d \in \Delta(X)$, $c \asymp d$.

Axiom VII-A2 (Reflexivity and Non-triviality)**.** $\preceq$ is reflexive and non-trivial.

Axiom VII-A3 (Stakes Transitivity)**.** For all $f, g, h, e, e' \in \mathcal{A}$, $\alpha, \beta \in (0, 1]$ such that $(f, h) \leqq (f_\alpha e, g_\alpha e)$ or $f(s) \sim g(s)$ for all $s \in S$, and $(f, h) \leqq (g_\beta e', h_\beta e')$ or $g(s) \sim h(s)$ for all $s \in S$, if $f_\alpha e \preceq g_\alpha e$ and $g_\beta e' \preceq h_\beta e'$, then $f \preceq h$.

Axiom VII-A4 (Pure Independence)**.** For all $f, g, h \in \mathcal{A}$ and for all $\alpha \in (0, 1)$ such that $f \asymp g$ and $f_\alpha h \asymp g_\alpha h$, $f \preceq g$ if and only if $f_\alpha h \preceq g_\alpha h$.

Axiom VII-A5 (Monotonicity)**.** For all $f, g \in \mathcal{A}$, if $f(s) \preceq g(s)$ for all $s \in S$, then $f \preceq g$.

Axiom VII-A6 (Consistency)**.** For all $f, g, h \in \mathcal{A}$ and $\alpha \in (0, 1)$ such that $(f_\alpha h, g_\alpha h) \leqq (f, g)$, if $f \asymp g$, then $f_\alpha h \asymp g_\alpha h$.

Axiom VII-A7 (Continuity)**.** For all $f, g, h \in \mathcal{A}$, the set $\{(\alpha, \beta) \in [0, 1]^2 \mid f_\alpha h \preceq g_\beta h\}$ is closed in $[0, 1]^2$.

Axiom VII-A8 (Continuity in Stakes)**.** For all $f, g, h \in \mathcal{A}$ with (f, g) not $\leqq$-maximal and $f(s), g(s) \succeq h(s)$ for all $s \in S$, $f \succeq g$ if and only if for all $\beta \in (0, 1)$, there exists $e \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $(f_\alpha e, (g_\beta h)_\alpha e) > (f, g)$ and $f_\alpha e \succeq (g_\beta h)_\alpha e$.

Axiom VII-A9 (Centering)**.** For all $f \in \mathcal{A}$ and $c \in \Delta(X)$, if $f_\alpha h \not\asymp c_\alpha h$ for all $h \in \mathcal{A}$ and $\alpha \in (0,1]$, then, for all $d \in \Delta(X)$ with $d \not\sim c$, there exists $h' \in \mathcal{A}$ and $\alpha' \in (0,1]$ such that $f_{\alpha'} h' \asymp d_{\alpha'} h'$.

Reflexivity and Non-triviality (VII-A2) and Monotonicity (VII-A5) are standard and require no further comment. Continuity (VII-A7) is a slight strengthening of the standard continuity axiom, and is related to axioms used elsewhere in the literature on incomplete preferences (see, for example, Dubra et al. (2004)). Indeed, in the presence of transitivity, independence and monotonicity, this axiom is equivalent to the standard one (see for example Gilboa et al. (2010, Lemma 3)). C-Completeness (VII-A1), which was introduced by Gilboa et al. (2010), simply says that preferences over constant acts are determinate. It translates the fact that the agent is assumed to be fully confident in her utilities; as stated in the Introduction, only confidence in beliefs is at issue here.

As concerns Stakes Transitivity (VII-A3), note firstly that transitivity in the case of incomplete preferences involves two distinct conditions: firstly, if $f \preceq g$ and $g \preceq h$, then one has determinate preferences between f and h; secondly, these preferences go in the appropriate direction – that is, $f \preceq h$. However, the former condition may be too strong. A decision maker may prefer spending \$10 on a bet on a certain ambiguous event to her current portfolio, no matter what her current portfolio is. Transitivity (applied repeated) implies that she prefers spending \$10 000 on 1000 bets on this same event to her current portfolio, whereas it does not seem unreasonable, under any of the standard interpretations of incompleteness cited in the Introduction, to have indeterminate preferences over these options.

Stakes transitivity weakens the first clause of the standard transitivity property, whilst retaining its second clause. More precisely, except for cases where f and g or g and h are extensionally equivalent, it demands determinate preference between f and h only when the decision maker's preferences between f and g and between g and h are determinate for stakes higher than the stakes involved in the choice between f and h.[10] It allows preferences to be indeterminate if this is not the case. In the example above, stakes transitivity thus allows indeterminacy of preferences concerning the \$10 000 bet, insofar as the stakes are higher than for a single \$10 bet. Importantly, in the presence of the other axioms, stakes

10. Recall from the discussion of richness in Section VII.2.2 that preferences at different stakes levels are given by preferences over appropriate mixtures of the acts in question.

transitivity implies that, whenever preferences are determinate, they go in the direction implied by the standard transitivity axiom. So, to the extent that one can speak of 'violations' of the standard axiom, they never result in preference cycles, but only in indeterminacy of preference where transitivity would have implied a determinate preference. [11] In this sense, this is a particularly mild weakening of transitivity. Note finally that, in the presence of VII-A4, stakes transitivity is equivalent to transitivity whenever preferences are complete.

A similar situation holds for Pure Independence (VII-A4). Whereas the standard independence axiom implies, firstly, that certain preferences are determinate, and secondly, that they go in a certain direction, pure independence simply states that whenever preferences are determinate, they go in the direction specified by the standard independence condition. Evidently it fully retains the intuitions behind the standard axiom, whilst accounting for the examples given in Section VII.1.1. Indeed, it can be thought of as an alternative way of extending the traditional independence axiom to the case of incomplete preferences, which separates the part of the standard axiom concerning determinacy of preference from the arguably more important part concerning direction of preference.

Consistency (VII-A6) is perhaps the most novel axiom and naturally so: it deals with the relationship between preferences at different stakes levels. It says that, if preferences between f and g are determinate, then preferences will be determinate between any mixtures of f and of g, as long as the stakes are not higher. In other words, if one has determinate preferences between two options at a given stakes level, then as the stakes fall, one retains the determinacy of the preferences. If the decision maker can choose between the options when there are hundreds of thousands of dollars at stake, then she can still choose when there are only tens of thousands at stake. As such, it is this axiom in particular which translates the idea that the higher the stakes, the more confidence is needed to take the choice. This is a fully behavioral axiom, which is in principle testable by, for example, comparing preferences at different stakes levels. (Of course, the other axioms are as behavioral as their standard counterparts.)

Of the final two axioms, Continuity in Stakes (VII-A8) is a largely technical axiom. The main direction states that, whenever f is preferred to g, then as the stakes in the choice are gradually increased (supposing they are not maximal), f may no longer be preferred

11. Perhaps it would be more appropriate to call these 'abstentions' from the transitivity axiom, reserving the term 'violation' for patterns of determinate preference that are incompatible with the axiom.

to g, but the most preferred act 'below' g (in the appropriate sense of being a mixture of g with an act dominated statewise by g and f) to which it is preferred will not suddenly 'jump' down with a slight increase in stakes. This basically ensures that the 'lower contour set' of an act changes gradually with an increase in the stakes involved in the decision. The axiom also includes the converse direction: if, for each act 'below' g there is a higher stakes level at which f is preferred to that act, then f is preferred to g. This direction is in fact implied by the other axioms whenever (f, g) is not $\leqq$-minimal, and thus can be dropped if one assumes, for example, that $\leqq$ has no non-trivial minimal elements.

The final axiom, Centering (VII-A9), is not required for the representation in general, but characterises the centeredness property of the confidence ranking (Section VII.2.3). It basically states that for any act and every constant act except at most one, if the stakes are low enough in the choice between them, the decision maker's preferences over them are determinate. The basic intuition is that, if there is sufficiently little at stake, the decision maker will hazard a preference in all these cases, albeit one in which she is not necessarily very confident.

VII.3.2 Result

The preceding axioms characterize the following representation of preferences.

Theorem VII.1. *Let $\preceq$ be a binary relation on $\mathcal{A}$. The following are equivalent.*

(i) $\preceq$ satisfies VII-A1-VII-A8.

(ii) There exists a nonconstant affine utility function $u : \Delta(X) \rightarrow \Re$, a continuous confidence ranking Ξ and a cautiousness coefficient $D : \mathcal{A} \times \mathcal{A} \rightarrow \Xi$ for Ξ such that, for all $f, g \in \mathcal{A}$, $f \preceq g$ if and only if

$$\text{(VII.1)} \qquad \sum_{s \in S} u(f(s)).p(s) \leqslant \sum_{s \in S} u(g(s)).p(s) \quad \text{for all } p \in D((f, g))$$

Furthermore, Ξ and D are unique, and u is unique up to positive affine transformation. Finally, under the conditions above, VII-A9 is satisfied if and only if Ξ is centred.

VII.3.3 Decisiveness and attitudes to choosing in the absence of confidence

We now undertake a basic comparative statics analysis of a decision maker's decisiveness under the model proposed above. Beyond giving a characterization of decisiveness in this model, the analysis will also allow us to verify the claim made in the Introduction that the cautiousness coefficient can be understood as capturing the decision maker's attitude to choosing in the absence of confidence.

Models of incomplete preferences admit a particularly simple comparison of decision makers, in terms of the completeness of their preference relations. For decision makers 1 and 2, if 2 weakly prefers f to g whenever 1 does, this is an indication that 2 is less prone to indeterminacy of preference than 1. Formally, we say that $\preceq^1$ is *less decisive* than $\preceq^2$ if $\preceq^1 \subseteq \preceq^2$.[12]

In order to characterise this relation in terms of the elements of the model, we require the relation $\sqsubseteq$ on families of sets, defined as follows. For two families of sets Ξ and Ξ', we write $\Xi \sqsubseteq \Xi'$ when, for every $\mathcal{C} \in \Xi$, there exists $\mathcal{C}' \in \Xi'$ with $\mathcal{C} \subseteq \mathcal{C}'$. We have the following result.

Proposition VII.1. *Let $\preceq^1$ and $\preceq^2$ satisfy axioms VII-A1–VII-A8, and be represented by (u_1, Ξ_1, D_1) and (u_2, Ξ_2, D_2) respectively. The following are equivalent:*

(i) *$\preceq^1$ is less decisive than $\preceq^2$.*

(ii) *u_2 is a positive affine transformation of u_1, $\Xi_2 \sqsubseteq \Xi_1$ and $D_2((f,g)) \subseteq D_1((f,g))$ for all $(f,g) \in \mathcal{A} \times \mathcal{A}$.*

Besides the utility functions, the main two elements in this result are the confidence ranking and the cautiousness coefficient. However, they may be understood as playing different roles. Consider two decision makers with the same utility function. Decision maker 1 has unanimity preferences à la Bewley: her confidence ranking contains a single set of probability measures $\mathcal{C}_1$ and the cautiousness coefficient sends all pairs of acts to that set. Decision maker 2, by contrast, has a rich confidence ranking Ξ_2, with a range of sets of probability measures, and an appropriate cautiousness coefficient, sending different pairs

12. Containment of the preference relations is equivalent to saying that, for all $f, g \in \mathcal{A}$, if $f \succeq^1 g$ then $f \succeq^2 g$.

of acts to different sets. As long as $\mathcal{C}' \subseteq \mathcal{C}_1$ for all $\mathcal{C}' \in \Xi_2$, 1 will be less decisive than 2. However, there seems to be something more precise to say about the relationship between the two decision makers. In particular, it appears that, on the one hand, 1 is less sensitive to the importance of decisions than 2 – if she prefers f over g at a given stakes level, then she has the same preference at any stakes level, no matter how high – but, on the other hand, 1 is confident of fewer beliefs than 2. In other words, there seems to be an aspect of belief (how confident one is of certain beliefs) as well as an aspect of taste (how willing one is to decide on the basis of beliefs in which one has a certain amount of confidence) mixed together in Proposition VII.1. To tease them apart, let us introduce the notion of confidence in preferences.

Definition VII.2. Let $\preceq$ satisfy axioms VII-A1–VII-A8. The *confidence-in-preferences* relation $\preccurlyeq$ on $\mathcal{A} \times \mathcal{A}$ is defined as follows: for any $f, g, f', g' \in \mathcal{A}$, $(f, g) \preccurlyeq (f', g')$ iff, for all $\alpha, \alpha' \in (0, 1], h, h' \in \mathcal{A}$ such that $(f_\alpha h, g_\alpha h) \equiv (f'_{\alpha'} h', g'_{\alpha'} h')$:

$$f_\alpha h \succeq g_\alpha h \Rightarrow f'_{\alpha'} h' \succeq g'_{\alpha'} h'$$

Definition VII.2 relies on the observation that, if a decision maker prefers f' to g' at a given stakes level but has indeterminate preferences between f and g at that level, then this can be taken as an indication that she is *more confident in her preference for f' over g' than in her preference for f over g.*[13] In other words, one can extract information about a decision maker's confidence in her preferences from the extent to which she holds specific preferences at given stakes levels. This is done according to the simple principle: the preferences that the decision maker holds at higher stakes are those in which she is more confident.

Given these considerations, we shall say that two decision makers are confidence equivalent if they have the same confidence-in-preferences relation.

Definition VII.3. Let $\preceq^1$ and $\preceq^2$ satisfy axioms VII-A1–VII-A8. $\preceq^1$ and $\preceq^2$ are *confidence equivalent* if $\preccurlyeq^1 = \preccurlyeq^2$.

Proposition VII.2. *Let $\preceq^1$ and $\preceq^2$ satisfy axioms VII-A1–VII-A8, and be represented by (u_1, Ξ_1, D_1) and (u_2, Ξ_2, D_2) respectively. $\preceq^1$ and $\preceq^2$ are confidence equivalent if and only if u_2 is a positive affine transformation of u_1 and $\Xi_1 = \Xi_2$.*

13. Recall from the discussion of richness in Section VII.2.2 that talk of preferences at different stakes levels is spelt out formally in terms of preferences over appropriate mixtures.

This proposition confirms that a decision maker's confidence in her preferences is entirely determined by her utilities and her confidence in beliefs. This is to be expected: to the extent that preferences are determined by utilities and beliefs, it is reasonable that confidence in preferences be determined by confidence in utilities – which is trivial in this model, because of the use of a single utility function – and confidence in beliefs, represented by the confidence ranking. The notion of confidence in preferences also helps shed light on the example above: decision makers 1 and 2 obviously have different confidence in preferences, and it is this difference, as much as any difference in attitude to choosing in the absence of confidence, that yields the difference in decisiveness. In fact, once differences in confidence in beliefs are accounted for, by comparing decision makers who have the same confidence in preferences, decisiveness is entirely determined by the cautiousness coefficient, as the following corollary of Propositions VII.1 and VII.2 shows.

Corollary VII.1. *Let $\preceq^1$ and $\preceq^2$ satisfy axioms VII-A1–VII-A8, be confidence equivalent, and be represented by (u, Ξ, D_1) and (u, Ξ, D_2) respectively. The following are equivalent:*

(i) $\preceq^1$ *is less decisive than* $\preceq^2$.

(ii) $D_2((f, g)) \subseteq D_1((f, g))$ *for all* $(f, g) \in \mathcal{A} \times \mathcal{A}$.

In summary, one can compare decision makers' decisiveness, that is, the pairs of acts for which they have determinate preferences. Comparison in terms of decisiveness corresponds to identical utilities (up to positive affine transformation) and appropriate relations of containment on the confidence rankings and cautiousness coefficients. Furthermore, one can define a decision maker's confidence in preferences on the basis of her preferences. Decision makers with the same confidence in preferences are precisely those who share the same utilities and confidence in beliefs, represented by a confidence ranking. Finally, for decision makers with the same confidence in preferences, differences in decisiveness are completely characterized by the relationship between their cautiousness coefficients. To the extent that such decision makers have the same confidence, differences in decisiveness must come down to differences in their attitudes to choosing on the basis of limited confidence. This supports the interpretation of the decision maker's cautiousness coefficient as capturing precisely her attitude to choosing in the absence of confidence.

VII.4 Incomplete preferences and choice

There may be situations in which indeterminate preferences have direct consequences in choice. Decision makers with indeterminate preferences may opt for the status quo, if it exists (Bewley, 2002); they may postpone the decision, if possible (Danan, 2003a; Kopylov, 2009); more generally, they may take a deferral option, if one is present (Hill, 2014). However, in many situations, such 'choice-avoidance mechanisms' are unavailable and the decision maker will have to make a choice, notwithstanding the incompleteness of her preferences. This essentially poses the question of how a decision maker with incomplete preferences 'completes' her preference relation in situations where she must choose.

This question is evidently relevant under all of the interpretations of incomplete preferences mentioned in the Introduction, though the form it takes sometimes depends on the interpretation adopted. For example, in the perspective proposed by Gilboa et al. (2010), the question of completion is that of the relationship between objectively and subjectively rational preferences. Note that a purely behavioral interpretation of the question can be given, by thinking of incompleteness in terms of deferral. The incomplete preference relation considered in previous sections can be thought of as representing the decision maker's behavior when a deferral option is available: when preferences are indeterminate, she defers. The question of completion is thus the question of how she would choose in situations where no deferral option is available. [14]

Consider a decision maker who is forced to choose between options over which she has indeterminate preference. Pre-theoretically, two sorts of strategies for deciding suggest themselves. One sort of strategy *respects confidence*: it uses only the beliefs that the decision maker holds with the level of confidence that corresponds to the stakes involved in the decision. The intuition is that, since these are the appropriate beliefs for decisions involving these stakes, they are the only ones she allows herself to rely on when deciding. Since they do not yield a determinate choice under representation (VII.1), the decision maker employs a different decision rule, involving, for instance, considerations of caution or an element of random choice. Another sort of strategy *goes on hunches*: it allows the decision maker to use all of her beliefs, irrespective of the confidence she has in them. The intuition here is

14. See Hill (2009) for further discussion of this interpretation of incompleteness, and Hill (2014) for a general treatment of deferral in the context of decision under uncertainty.

that, whilst the decision maker would not form preferences on the basis of a hunch (a belief in which she has insufficient confidence) – or decide on a hunch if she could avoid it by deferring the decision, for instance – when she is forced to decide she may as well mobilize all of her beliefs – even hunches. Given that she is relying on more beliefs, the decision maker may be able to decide using the 'unanimity' rule employed in representation (VII.1); if not, she will require a different decision rule.

We would suggest that this distinction corresponds to an important difference between possible reactions to forced choice, under all of the interpretations cited above. It seems that some decision makers in certain situations – an entrepreneur following her 'gut feelings', or a general going on his intuition in the heat of a battle, for example – rely on beliefs in which they have insufficient confidence when called on to decide, whereas others in other situations – a governor deciding whether to permit the construction of a nuclear plant, or a doctor deciding on treatment for a patient, for example – only limit themselves to beliefs that they hold with sufficient confidence given the decision in hand. In the perspective proposed by Gilboa et al. (2010), a strategy that respects confidence only uses the beliefs that are 'objectively defendable' to form subjectively rational preferences, and in this sense is close to the representation given in that paper. However, it does not seem unreasonable, when forming one's (personal) subjectively rational preferences, to rely on beliefs that one cannot convince others of with sufficiently strong arguments: that is, to adopt a strategy that goes on hunches. Finally, the distinction takes a particularly simple form under the group interpretation of incompleteness: in cases of disagreement among the group, a strategy that respects confidence forms preferences accounting for the full scope of the disagreement, whereas a strategy that goes on hunches chooses what the board of directors deem preferable, ignoring the others' opinions. Whilst some groups, such as certain associations, may sometimes use the former strategy, others, for example many firms, often seem to use the latter one.

To tackle the question of choice on the basis of incomplete preferences, we augment the framework introduced in the previous sections with a binary relation, $\preceq^c$, over the set of acts. It represents the 'completion' of the decision maker's preference $\preceq$ from Sections VII.2 and VII.3. The issue of choice on the basis of incomplete preferences can be tackled by considering axioms on these preference relations and the relationship between them, such as the following.

Axiom VII-B1 (Forced Choice)**.** $\preceq^c$ is complete.

Axiom VII-B2 (Benchmark on Certainty)**.** For all $f, g \in \mathcal{A}$, $g \prec^c f$ if and only if there exist $c, d \in \Delta(X)$ with $c \succ d$ such that $f_\alpha h \succeq c_\alpha h$ for some $h \in \mathcal{A}$ and $\alpha \in (0,1]$ and $g_{\alpha'} h' \not\succeq d_{\alpha'} h'$ for all $h' \in \mathcal{A}$ and $\alpha' \in (0,1]$.

Axiom VII-B2^S (Stakes Benchmark on Certainty)**.** For all $f, g \in \mathcal{A}$, $g \prec^c f$ if and only if there exist $c, d \in \Delta(X)$ with $c \succ d$ such that $f_\alpha h \succeq c_\alpha h$ for some $h \in \mathcal{A}$ and $\alpha \in (0,1]$ with $(f_\alpha h, c_\alpha h) \geqq (f, g)$, and $g_{\alpha'} h' \not\succeq d_{\alpha'} h'$ for all $h' \in \mathcal{A}$ and $\alpha' \in (0,1]$ with $(g_{\alpha'} h', d_{\alpha'} h') \geqq (f, g)$.

Forced Choice (VII-B1) states that the completed preferences are indeed complete; it translates the fact that the decision maker must choose. The other two axioms, Benchmark on Certainty (VII-B2) and Stakes Benchmark on Certainty (VII-B2^S), involve the same basic intuition. Incomplete preferences provide a crude indication of the relative worth of different acts for the decision maker; one way of getting a more refined judgement is by considering how the acts compare to constant acts. Thus, even if the decision maker's preferences between two acts, say f and g, are indeterminate, she may have determinate preference for f over a particular constant act (eg. a sure \$5), whilst not having a determinate preference for g over an inferior constant act (eg. a sure \$4). The axioms both demand that in precisely these cases she strictly prefers f to g when forced to choose. In other words, an act is chosen over another in situations of forced choice precisely when it fairs favourably with respect to the other act in the comparison with constant acts according the initial (incomplete) preference relation. Both VII-B2 and VII-B2^S characterise a form of conservative decision making, insofar as all that counts in determining the preference completion are the constant acts that the acts in question are preferred to. The difference between them is that whilst Benchmark on Certainty (VII-B2) licences all comparisons with constant acts, no matter the stakes level, Stakes Benchmark on Certainty (VII-B2^S) only involves stakes-corrected comparisons: comparisons where all the pairs of acts considered involve sufficiently high stakes. As the following results show, these two axioms characterize behaviorally the difference between the two strategies described above.

Theorem VII.2. *Let $\preceq$ satisfy VII-A1–VII-A8, and be represented according to (VII.1) by (u, Ξ, D). Then* *(i)* $\preceq$ *and* $\preceq^c$ *satisfy VII-B1 and VII-B2 if and only if for all* $f, g \in \mathcal{A}$,

$$\text{(VII.2)} \qquad f \preceq^c g \ \textit{ iff } \ \min_{p \in \bigcap_{C \in \Xi} C} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in \bigcap_{C \in \Xi} C} \sum_{s \in S} u(g(s)).p(s)$$

(ii) $\preceq$ *and* $\preceq^c$ *satisfy VII-B1 and VII-B2*^{S} *if and only if for all* $f, g \in \mathcal{A}$,

$$\text{(VII.3)} \qquad f \preceq^c g \ \textit{ iff } \ \min_{p \in D((f,g))} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D((f,g))} \sum_{s \in S} u(g(s)).p(s)$$

Both Benchmark to Certainty (VII-B2) and Stakes Benchmark to Certainty (VII-B2^S) characterize cautious decision making, insofar as decision makers who satisfy them choose on the basis of the minimum expected utility taken over a set of probability measures (Gilboa and Schmeidler, 1989). They differ, however, in the set over which the minimum is taken. Stakes Benchmark to Certainty (VII-B2^S) implies that the decision maker uses the set of probability measures corresponding to the level of confidence appropriate for the stakes involved in the choice. This set represents the beliefs that she holds with sufficient confidence given the stakes. So, although she may not be able to form a determinate preference on the basis of these beliefs under the representation proposed in Section VII.3, she continues to use them when she is forced to choose: she chooses the act with the highest worst-case expected utility, given the beliefs in which she has sufficient confidence. By contrast, Benchmark to Certainty (VII-B2) yields a representation involving the smallest set of probability measures in her confidence ranking. This set encapsulates all her beliefs, even those in which she has little confidence. So, while the decision maker uses only beliefs held with sufficient confidence to form preferences according to representation (VII.1), she does allow herself to rely on beliefs in which she has insufficient confidence when forced to choose: she chooses the act that has the highest worst-case expected utility, given all of her beliefs.

Obviously the decision makers satisfying Stakes Benchmark to Certainty employ a cautious strategy that respects confidence to decide when they are forced to choose, whereas those satisfying Benchmark to Certainty employ a cautious strategy that goes on hunches. Hence, between cautious strategies for ‘completing’ incomplete preferences, the difference between VII-B2^S and VII-B2 characterizes precisely the difference between the strategy that respects confidence and the one that goes on hunches. The ability to capture in a simple and precise way both of these pre-theoretically conceivable, and apparently relevant, strategies for deciding when one is not sure could be considered as a strength of the present approach, and in particular of the use of the notion of confidence ranking. To the knowledge of the author, this is the first model in the literature capable of capturing this distinction.

Whilst representations (VII.2) and (VII.3) involve the maxmin EU rule, other strategies for deciding in situations of forced choice may involve different decision rules, such as the maxmax EU rule, the α-maxmin EU rule (Ghirardato et al., 2004) or a random decision rule in the style of Gul and Pesendorfer (2006). There are two versions of each of these strategies: one that respects confidence – using the representation (VII.3) with the appropriate decision rule in place of the maxmin EU one – and one that goes on hunches – where the decision rule is used in a version of representation (VII.2). So the difference between strategies respecting confidence and those that go on hunches is orthogonal to the issue of which decision rule is used in the case of forced choice. To this extent, the representations studied above can be thought of as examples drawn from a general family of models of forced choice on the basis of incomplete preferences.

Finally, let us comment on the case of a decision maker with a centered confidence ranking. Faced with a decision in which she is forced to choose, but where there is low confidence in the relevant beliefs, if such a decision maker goes on hunches, then she acts precisely like a subjective expected utility maximizer. In this case, the distinction between the two strategies for preference completion may be thought of as offering a new perspective on the debate between Bayesians and non-Bayesians: to the question of whether decision makers can form precise probabilities (raised, for example, by Gilboa et al. (2009)), it adds the question of whether they should choose on the basis of them even if they could form them.

VII.5 Confidence and indeterminacy in markets

As a further exploration of the wider implications of the model, we briefly consider some consequences for risk sharing in financial markets. Recall that one interpretation of indeterminacy of preferences is in terms of status quo choice, if there is a status quo option.

A status quo is present in a market setting: it is simply the option of not trading. Bewley (1989) and Rigotti and Shannon (2005) have considered the consequences of the unanimity model (à la Bewley) in a market setting, interpreting indeterminacy of preferences as the choice of not trading. We do the same here for preferences represented according to (VII.1).

We consider a standard Arrow-Debreu exchange economy with a complete set of (non-negative) state-contingent commodities on a finite state space S. The set of acts $\mathcal{A}$ is defined as in Section VII.2.1, with the set of outcomes specified by $X = \Re_+$. A state-contingent commodity is a vector in $\Re^S_+$, and can be naturally assimilated with the corresponding element in $\mathcal{A}$.[15] With slight abuse of notation, a constant state-contingent commodity yielding outcome w in every state will be denoted w. The economy has finitely many agents, indexed by $i = 1 \ldots n$. Each has preferences $\preceq^i$ over $\mathcal{A}$ (and hence over $\Re^S_+$), represented as in (VII.1) for a stakes relation $\leqq^i$. Each agent thus has a utility function $u^i : \Re_+ \to \Re$, a confidence ranking Ξ^i on S and a cautiousness coefficient D^i. We assume that all u^i are differentiable, strictly concave and strictly increasing. Note that, since expected utility preferences and unanimity preferences à la Bewley are special cases of (VII.1), the economy may contain agents with these sorts of preferences. The aggregate endowment is $e \in \Re^S_{++}$. Finally, an allocation $(x^1, \ldots, x^n) \in (\Re^S_+)^n$ is said to be *feasible* if $\sum_i x^i = e$, it is *interior* if $x^i_s > 0$ for all i and s, and it is a *full insurance allocation* if all the x^i are constant.

Definition VII.4. An allocation $(y^1, \ldots, y^n)$ *Pareto dominates* the allocation $(x^1, \ldots, x^n)$ if, for each agent i, either $y^i \succ^i x^i$ or $y^i = x^i$.

A feasible allocation $(x^1, \ldots, x^n)$ is *Pareto optimal* if there is no feasible allocation that Pareto dominates it.

This notion of Pareto optimality is very close to that studied by Fon and Otani (1979). The notion of Pareto dominance employed says that an allocation dominates another ex-

15. State-contingent commodities correspond to acts whose consequences are degenerate lotteries; $\Re^S_+$ thus corresponds to a proper subset of $\mathcal{A}$. Nevertheless, for each agent, thanks to the continuity of her utility function, every lottery over $\Re_+$ (ie. element of $\Delta(X)$) has a certainty equivalent in $\Re_+$, so, given her utility function, her preferences over $\mathcal{A}$ are completely determined by her preferences over $\Re^S_+$. Hence, although we assume preferences over $\mathcal{A}$, this is equivalent in this setup to assuming preferences over $\Re^S_+$; similarly, properties of preferences can be formulated either in terms of $\mathcal{A}$ or $\Re^S_+$. We continue to use the notation introduced in Section VII.2, and in particular the generic symbols $f, g, h \ldots$ for acts; we use standard vector notation and generic symbols $x, y, z \ldots$ for commodities.

actly when all agents who trade contingent commodities strictly prefer their new commodity to their old one. This is a natural notion in the context of incomplete preferences where indeterminacy is interpreted in terms of sticking to the status quo: it supposes that agents who do not have strict preference for trade – either because they consider the commodity on offer not to be better than what they have, or because they do not have sufficient confidence to form determinate preferences – stick to their initial endowment.

For technical reasons, we make several assumptions on the notions of stakes and on the preferences. Since they are tangential to the general discussion in this section, they are simply noted here; statement of the relevant properties and discussion are relegated to Appendix VII.B.

Assumption VII.1. For each $1 \leqslant i \leqslant n$, the stakes relation $\leqq^i$ is monotone decreasing, and the preference relation $\preceq^i$ satisfies indifference consistency and full support.

Under this assumption, we have a characterisation of Pareto optima. For a contingent commodity $x \in \Re^S_+$ and an agent i, let [16]

$$\Pi^i(x) = \left\{ \left(\frac{p(s_1)u^{i\prime}(x_{s_1})}{\sum_{t\in S} p(t)u^{i\prime}(x_t)}, \ldots, \frac{p(s_{|S|})u^{i\prime}(x_{s_{|S|}})}{\sum_{t\in S} p(t)u^{i\prime}(x_t)} \right) \mid p \in \bigcap_{z \neq x} ri(D^i(x,z)) \right\}$$

Theorem VII.3. *Under Assumption VII.1, an interior allocation $(x^1, \ldots, x^n)$ is Pareto optimal iff $\bigcap_i \Pi^i(x^i) \neq \varnothing$.*

The intuition behind this result is analogous to similar results in the literature (Rigotti and Shannon, 2005; Rigotti et al., 2008): $\Pi^i(x)$ is the set of supports of the strict upper contour set of x under $\preceq^i$, and an allocation is Pareto optimal if and only if the intersection of all such sets is non-empty. The theorem has several immediate consequences.

Corollary VII.2. *Let Assumption VII.1 hold, and suppose that the aggregate endowment is constant across states.*

(i) An interior full insurance allocation $(x^1, \ldots, x^n)$ is Pareto optimal iff $\bigcap_i \bigcap_{z \neq x^i} ri(D^i(x^i, z)) \neq \varnothing$.

16. For a set A, $ri(A)$ is the relative interior of A. Note that, if A is a singleton, $ri(A) = A$.

(ii) If, for each i, there exists $\mathcal{C}^i \subseteq \Delta(S)$ with $\Xi^i = \{\mathcal{C}^i\}$, then there exists an interior full insurance Pareto optimal allocation iff $\bigcap_i ri(\mathcal{C}^i) \neq \emptyset$. In this case, every full insurance allocation is Pareto optimal.

The first corollary, which is a simple consequence of the fact that $\Pi^i(x^i) = \bigcap_{z \neq x^i} ri(D^i(x^i, z))$ when x^i is constant, is a general characterisation of Pareto optimality of an interior full insurance allocation under representation (VII.1). It is in the style of existing results, such as Billot et al. (2000); Rigotti and Shannon (2005); Rigotti et al. (2008). Unlike these cases, the existence of a full insurance Pareto optimal allocation does not imply that all full insurance allocations are Pareto optimal, because, in general, the relevant sets of probability measures may differ depending on the constant commodity.

The second corollary involves the special case of representation (VII.1) where the confidence ranking is degenerate: this is essentially the unanimity model of preferences à la Bewley. The result differs slightly from that of Rigotti and Shannon (2005, Corollary 2), which also concerns the Bewley model, insofar as their result involves the intersection of the sets of probability measures of the different agents, whereas ours uses the intersections of their relative interiors. This difference is due to the fact that they take the strict preference relation as primitive, use a slightly different representation from that used here and take a stricter notion of Pareto optimality. [17]

Comparing the two corollaries brings out the relationship between the consequences of representation (VII.1) for risk sharing in markets and those of the unanimity model à la Bewley. Concerning the question of whether an interior full insurance allocation is Pareto optimal or not, an economy with agents represented by (VII.1) is roughly equivalent to an economy where each agent is replaced by an agent with unanimity preferences à la Bewley who takes as her set of probability measures the intersection of all the sets in her confidence ranking that are relevant for choices involving the commodity she is allocated. [18] *Grosso modo*, if each agent in the economy for whom more confidence is required to take decisions with higher stakes simply ignored the stakes, and always chose as if the stakes were at the lowest possible level for the commodity she is allocated, this would make little difference to whether the allocation is Pareto optimal or not. Although it does not follow that the set of Pareto optima would be the same as for an economy containing only agents with

17. They use the representation axiomatised by Bewley (2002); see Section VII.1.3.

18. This is a rough statement because, for a family $\{\mathcal{C}_i |\ i \in I\}$ of closed sets, it is not necessarily the case that $\bigcap_{i \in I} ri(\mathcal{C}_i) = ri \bigcap_{i \in I} \mathcal{C}_i$.

preferences à la Bewley, because the relevant sets of probability measures may depend on the allocation, this highlights some similarities between economies with agents à la Bewley and economies with agents represented according to (VII.1).

Things are considerably different, however, regarding the question of how fast Pareto optima can be reached. (To the extent that, as noted in the proof of Theorem VII.3, Pareto optima correspond to appropriately defined equilibria, this is closely related to the question of how fast the economy can arrive at equilibrium.) [19] There is often a simple, if idealised, fastest way to achieve a Pareto optimum. In particular, whenever a non-Pareto optimal allocation $(x^1, \ldots, x^n)$ is dominated by a (feasible) Pareto optimal one $(y^1, \ldots, y^n)$, then there is a 'one-step' move to a Pareto optimum, which is acceptable to all agents – namely, each agent swaps x^i for y^i. (This set of 'swaps' corresponds to a set of simultaneous trades between the agents.) Whenever this is the case, we say that $(y^1, \ldots, y^n)$ is *one-step accessible* from $(x^1, \ldots, x^n)$. In economies where agents have expected utility preferences, preferences à la Bewley or preferences represented by many of the standard non-expected utility theories proposed in the literature (and in particular those considered by Rigotti et al. (2008)), any Pareto dominated allocation is Pareto dominated by a Pareto optimum. In other words, any allocation has a Pareto optimum that is one-step accessible from it. This is not necessarily true for economies containing agents represented by (VII.1), as the following example shows.

Consider a two-agent economy with two states of the world, s_1 and s_2, and suppose that each agent i has constant relative risk aversion γ^i (so the utility function is $u^i(x) = \frac{x^{1-\gamma^i}}{1-\gamma^i}$ if $\gamma^i \neq 1$ and $u^i(x) = \ln x$ if $\gamma^i = 1$). Agent 1's preferences are represented by (VII.1), where the stakes in the choice between x and y are given by $\max_s |x(s) - y(s)|$. [20] She has the following centred confidence ranking: $\{\{p \in \Delta(\Sigma) |\ 0.5 - \epsilon \leqslant p(s_1) \leqslant 0.5 + \epsilon\} |\ \epsilon \in [0, 0.45]\}$. Note that since each set in the confidence ranking is uniquely specified by an $\epsilon \in [0, 0.45]$, the cautiousness coefficient is entirely specified by a function from pairs

19. Rigotti and Shannon (2005) address a related question with their notion of 'equilibrium with inertia', which, approximately, is an equilibrium which Pareto dominates the initial endowment. The example below shows that, by contrast with economies whose agents have preferences à la Bewley, equilibria with inertia may not exist, even if equilibria do exist, in economies whose agents have preferences represented by (VII.1).

20. As noted in footnote 15, although the stakes relation is defined on contingent commodities, this yields a well-defined stakes relation on acts. It is straightforward to check that this relation satisfies the properties assumed in this paper (see Section VII.2.2 and Appendix VII.A) and in Assumption VII.1.

of acts to values of ϵ. Using this formulation, the cautiousness coefficient is given by $D^1((x,y)) = \min\{\eta \max_s |x(s) - y(s)|, 0.45\}$ for $\eta > 0$, where η characterises the agent's attitude to choosing in the absence of confidence (see Section VII.3.3). Agent 2 is an expected utility decision maker with probability measure assigning 0.5 to both states.

We consider a case with no aggregate risk in the economy: the sum of allocations is w in both states. Hence allocations are of the form $((\delta_1 w, \delta_2 w), ((1-\delta_1)w, (1-\delta_2)w))$ for $\delta_1, \delta_2 \in [0,1]$. It is easy to check, given Theorem VII.3, that the only Pareto optima are full insurance allocations. Now consider the risky endowment $(x^1, x^2) = ((\delta w, (1-\delta)w), ((1-\delta)w, \delta w))$, where $\delta \in (\frac{1}{2}, 1]$. It would seem that a natural 'one-move' trade yielding a Pareto optimal allocation would be for 2 to give 1 $((\frac{1}{2} - \delta)w, (\delta - \frac{1}{2})w)$. It is easy to check that $\frac{1}{2}w \succ^2 x_2$. Moreover, $\sum_s p(s)u^1(\frac{1}{2}w) > \sum_s p(s)u^1(x^1_s)$ for all $p \in \bigcap_{x' \neq x^1} ri(D^1((x^1, x')))$. Were the agents to ignore the stakes and always choose as if the stakes were at their lowest level, this would be sufficient for the trade to be acceptable to both agents: that is, for the full insurance allocation $(\frac{1}{2}w, \frac{1}{2}w)$ to be one-step accessible from (x^1, x^2). However, if they take the stakes into account as specified by representation (VII.1), there is a stronger condition that is required for the trade to be acceptable to agent 1, namely that $\sum_s p(s)u^1(\frac{1}{2}w) \geqslant \sum_s p(s)u^1(x^1_s)$ for all $p \in D^1((x^1, \frac{1}{2}w))$, with strict inequality for some p. By straightforward calculation, this condition holds if and only if [21]

$$p \leqslant \frac{\frac{1}{2}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}} \qquad \text{for all } p \in D^1((x^1, \frac{1}{2}w)) \tag{VII.4}$$

Hence, by the definition of D^1, $\frac{1}{2}w \not\succ^1 x^1$ whenever

$$\min\{\eta w(\delta - \frac{1}{2}), 0.45\} + 0.5 > \frac{\frac{1}{2}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}} \tag{VII.5}$$

This inequality has solutions for various values of the parameters: it is straightforward to check, for example, that when $\gamma^1 = 2$, $\delta = \frac{3}{4}$, $w = 1500$, $\eta = 0.001$, the inequality is satisfied and so $\frac{1}{2}w \not\succ^1 x^1$. In such cases, the Pareto optimum $(\frac{1}{2}w, \frac{1}{2}w)$ is not one-step accessible from (x^1, x^2). In fact, by a similar argument, one can show that there are cases where no Pareto optimal allocation is one-step accessible from (x^1, x^2).

21. Here we consider the case where $\gamma_1 \neq 1$; the case of $\gamma_1 = 1$ can be treated similarly.

Proposition VII.3. *There may exist Pareto dominated allocations from which no Pareto optimal allocation is one-step accessible.*

This phenomenon is basically a consequence of the dependence on stakes in representation (VII.1), which allows agents to have determinate preferences at low stakes levels that they may withdraw at higher stakes levels. Whereas it is the former preferences – and in particular the probability measures corresponding to low levels of stakes – which determine whether an allocation is Pareto optimal or not, the latter preferences – and the associated larger sets of probability measures – determine whether an agent accepts a given trade or not. If all agents were indifferent to the stakes, and formed preferences using the minimal sets of probability measures in their confidence rankings (that is, as if the stakes were at their lowest level), then any Pareto dominated allocation would indeed be Pareto dominated by a Pareto optimal one. However, whenever there is an agent who takes the stakes into account according to representation (VII.1), she may not be confident enough in her strict preference for that Pareto optimal allocation over her initial endowment to choose the former at the appropriate level of stakes, and so sticks to the status quo. She refrains from trading, and the 'one-step' move to the Pareto optimum is blocked.

We have already mentioned one interpretation of this result in terms of maximal speed of convergence to equilibrium. It indicates a non-trivial bound on how fast a Pareto optimal allocation can be reached: allowing any conceivable way of constructing a set of simultaneous trades (as unfeasible as it may be in practice), it may still be impossible to get to a Pareto optimum by a single set of trades if the market contains agents who incorporate confidence into their preferences, and who do not trade when they lack sufficient confidence. Another interpretation is in terms of the restrictions placed on the (theoretical) power of a social planner. In standard general equilibrium models, as well as the market under uncertainty models mentioned above, a suitably intelligent social planner who knows the agents' preferences could propose a set of simultaneous trades that would be accepted by all agents and that would bring the market to a Pareto optimum. This relies on the fact that, in these models, for each allocation, there is a Pareto optimum that is one-step accessible. That this is not necessarily the case in the current model attests to the limited influence of such a social planner: even if he had all the information about preferences (and infinite computational power), the social planner might not be able to propose a set of simultaneous trades that leaves the economy in a Pareto optimum and is acceptable to all. The agents' tendency to demand more confidence in beliefs when the stakes are higher mean that he may not be able to persuade some of them to shift from the endowment to a Pareto optimal allocation when the stakes involved in the change are high, though they would have accepted the trade if the stakes were low. Confidence, combined with the status quo interpretation of indeterminacy of preference, can hinder Pareto enhancing intervention in the market.

The natural question is, of course: how fast can a Pareto optimum be reached? Put in terms of the second interpretation offered above, this amounts to asking how many times a social planner has to intervene to bring the economy to a Pareto optimum. Let us say that a feasible allocation $(y^1, \ldots, y^n)$ is *m-step accessible* from $(x^1, \ldots, x^n)$ if there is a sequence of $m-1$ feasible allocations, the first of which Pareto dominates $(x^1, \ldots, x^n)$, the last of which is Pareto dominated by $(y^1, \ldots, y^n)$, and each of which is Pareto dominated by its successor. A Pareto optimum which is not one-step accessible may be m-step accessible: this means that, under ideal conditions, it can be reached not with a single set of trades that is acceptable to all, but rather after m consecutive sets of trades, each of which is acceptable to all. If, for a given allocation, there is a Pareto optimum that is m-step accessible and none that is m'-step accessible for $m' < m$, this can be thought of as a lower bound on the how fast the economy can come to a Pareto optimum: it requires at least m sets of simultaneous trades. A social planner has to intervene at least m times. There are, however, allocations from which no Pareto optimum is accessible in a finite number of steps.

Proposition VII.4. *There may exist Pareto dominated allocations from which no Pareto optimal allocation is m-step accessible, for any finite m.*

When there are agents whose preferences incorporate their confidence in beliefs, and who stick to the status quo when they do not have enough confidence to take a choice, it may thus be theoretically impossible for the market to arrive at a Pareto optimal allocation in finite time. Because, quite simply, there may not exist a finite sequence of sets of trades, where all agents have sufficient confidence to accept the trades and where the sequence reaches a Pareto optimum. Confidence, combined with taking the status quo option – and not trading – when one is not sufficiently confident in any option, adds considerable friction into the economy.

VII.6 Conclusion

Decision makers may have incomplete preferences. Moreover, they may be more or less confident in their beliefs. In this paper, a theory which relates incompleteness of preferences to confidence in beliefs was proposed. It is based on the following maxim: one has a determinate preference over a pair of acts if and only if one's confidence in the beliefs needed to form the preference matches up to the stakes involved in the choice between the acts. In the absence of sufficient confidence, preferences are indeterminate.

A formal decision rule conforming to this maxim was proposed. The decision maker's confidence in her beliefs is modelled by a confidence ranking – a nested family of sets of probability measures. A cautiousness coefficient assigns to any decision a level of confidence relevant for that decision (represented formally by a set in the confidence ranking), which is determined by the stakes involved. The decision rule according to which one act is preferred to another if it has higher expected utility according to all the probability measures in the appropriate set was axiomatised. Moreover, comparative statics analysis of the relative decisiveness of decision makers, as well as of their confidence in preferences, was undertaken, and the cautiousness coefficient was seen to correspond to the decision maker's attitude to choosing in the absence of confidence.

It was argued that the choice-theoretic properties that distinguish the proposed model from the standard Bewley model of incomplete preferences are both axiomatically mild (they involve allowing indeterminacy where the Bewley model demands determinate preference), and behaviorally reasonable, under most of the existing interpretations of incomplete preferences. Moreover, the question of the 'completion' of incomplete preferences – which is relevant under all of the aforementioned interpretations, in particular to handle situations where a choice is required – was considered. The introduction of the notion of confidence allows the identification of two strategies for preference completion. One strategy respects confidence, insofar as it only relies on the beliefs that the decision maker holds to the appropriate level of confidence given the stakes involved in the decision. The other strategy goes on hunches, to the extent that it mobilizes all of the decision maker's beliefs, even those in which she has little confidence, in situations where she is forced to decide. It was argued that each of these strategies may be pertinent in different decision situations under the various interpretations of incomplete preferences, and an axiomatic characterization of the two strategies was proposed.

Finally, possible consequences of the model in a market setting were considered, where indeterminacy of preferences is interpreted in terms of refusal to trade. On the one hand, a characterisation of Pareto optima was provided, which is analogous to similar results for other decision models. On the other hand, it was shown that, unlike other models, there may exist Pareto dominated allocations that are not dominated by any Pareto optimum. This indicates that the incorporation of confidence can add a considerable friction to the economy: it may be theoretically impossible for the market to come to a Pareto optimum by a single set of trades accepted by all. Moreover, there are cases where no finite sequence of sets of trades, accepted by all, can bring the market to a Pareto optimum.

VII.A Notions of stakes

There are many notions of stakes, each of which generates a particular stakes relation. The weaker the properties assumed on the stakes relations, the more notions of stakes satisfy them, and the wider any results involving them apply. In this appendix, we briefly discuss the weakness of the properties assumed in this paper, and hence the generality of the results in Sections VII.3 and VII.4, by considering some possible notions of stakes. The following natural notions of stakes all yield stakes relations satisfying the five properties stated in Section VII.2.2: [22] the stakes in the choice between f and g are given by

(i) the maximum of the negation of the utility of the least preferred consequence which could be obtained, taken over f or g

(ii) the maximum utility of the most preferred consequence which could be obtained by f or g

(iii) the probability that either of f or g takes a value below some threshold, calculated using a given probability measure (possibly taken from the decision maker's confidence ranking).

Now consider the following notions of stakes: the stakes in the choice between f and g are given by

(iv) the difference between the utility of the least preferred consequence which could be obtained by f and the utility of the least preferred consequence which could be obtained by g

(v) the maximum absolute value of the difference between the utility of $f(s)$ and the utility of $g(s)$, taken over $s \in S$

(vi) the difference between the expected utilities of f and g, calculated using a given probability measure (possibly taken from the decision maker's confidence ranking).

It is straightforward to check that these notions generate stakes relations satisfying all of the properties in Section VII.2.2 except for richness. However, they do satisfy the following weaker richness condition:

22. Richness is satisfied by these notions of stakes under the assumption that the utility function is unbounded above; the weakening of richness discussed below is satisfied in the absence of this assumption.

(Richness') For all $f, f', g, g' \in \mathcal{A}$ such that $f(s) \nsim g(s)$ for some $s \in S$ and (f', g') is not $\leqq$-minimal, there exists $(f'', g''), (f''', g''') \in \|(f, g)\|$ such that $(f'', g'') \leqq (f', g') \leqq (f''', g''')$.

where $\|(f, g)\|$ is the set of pairs of acts (f', g') which can either be obtained from f and g by mixing them in the same proportion with another act, or such that f and g can be obtained from f' and g' by such a mixture. [23] Rather than demanding that every (non-trivial) pair of acts can be shifted as far up or down the stakes ordering as desired by mixing with a third act, richness' simply asks that, for every (non-trivial) pair of acts and (non-minimal) stakes level, there is a pair of acts below (respectively, above) that level, such that either the latter pair can be obtained from the former pair by mixing, or the former can be obtained from the latter by mixing. It evidently retains the intuition behind the original richness property – in particular, the pairs in $\|(f, g)\|$ correspond to versions of the choice between f and g but with potentially different stakes. Moreover, with a corresponding modification to the axioms in Section VII.3, one can prove a result analogous to (but stronger than) Theorem VII.1 for stakes relations satisfying the first four properties in Section VII.2.2 and richness'. We have chosen to work with richness in the bulk of this paper to avoid burdening the reader with notation.

Note moreover that the theory proposed by Nau (1992) is essentially a member of the class of representations of the form (VII.1) which takes as stakes notion (v). Hence the result just mentioned contains Nau's as a special case.

Finally we remark that defining the stakes relation on pairs of acts as done in this paper is the simplest, but far from the only possibility. It could also have been defined on triples in $\mathcal{A} \times \mathcal{A} \times \Gamma$, where Γ can have several interpretations. For example, Γ can be understood as a set of context indices; hence dependence of the stakes on the context can accommodated. Alternatively, Γ could be interpreted as the status quo, if there is one; using this relation, one can capture dependence of the stakes on the status quo. It is straightforward to adapt the properties given in Section VII.2.2 to such stakes relations; corresponding modifications to the axioms in Section VII.3 yield similar results where the stakes may depend on factors other than the two acts on offer.

23. The formal definition is as follows: for all $f', g' \in \mathcal{A}$, $(f', g') \in \|(f, g)\|$ if and only if there exists $\alpha \in (0, 1]$, $h \in \mathcal{A}$ such that $f'(s) \sim (f_\alpha h)(s)$ and $g'(s) \sim (g_\alpha h)(s)$ for all $s \in S$, or $f(s) \sim (f'_\alpha h)(s)$ and $g(s) \sim (g'_\alpha h)(s)$ for all $s \in S$.

VII.B Properties involved in Assumption VII.1

Readers familiar with the literature on general equilibria in the absence of completeness or transitivity might expect these results to be directly applicable to the case considered in Section VII.5. This is in fact not straightforwardly possible in general, for two reasons. Firstly, the weakening of transitivity in representation (VII.1) implies that, for certain notions of stakes, preferences represented by (VII.1) may not be convex.[24] Secondly, since the weak preference order is taken as primitive, it does not follow from the axioms in Section VII.3.1 that the strict preference order is continuous.[25]

Assumption VII.1 deals with these issues, by assuming the following properties of the stakes relation and of the preferences.

Monotone decreasing For all $f, g \in \mathcal{A}$ and $\alpha, \beta \in [0,1]$, if $\alpha \leqslant \beta$, then $(f, g_\alpha f) \leqq (f, g_\beta f)$.

Indifference consistency For all $f, g, h \in \mathcal{A}$ and $\alpha \in (0,1)$ such that $f \asymp g$ and $f_\alpha h \asymp g_\alpha h$, if $f \sim g$, then $f_\alpha h \sim g_\alpha h$.

Full support For all $s \in S$, there exists $c_s \in \Re_{++}$ such that, for all $h \in \mathcal{A}$ and $\alpha \in (0,1]$, $(1_s)_\alpha h \succ (c_s)_\alpha h$.[26]

The monotone decreasing property of stakes states that as one considers choices between acts that are 'closer' to each other (in the sense of mixtures), the stakes decrease. This property is satisfied by several of the notions of stakes mentioned in Appendix VII.A: for example, stakes as the minimum utility that could be obtained from either of the acts, as the difference between the minimum utilities yielded by the two acts, or as the maximum utility difference between the two acts (taken over the set of states).

24. Consider a stakes relation where $(f, g_\alpha h) > (f, g), (f, h)$, for some $\alpha \in (0,1)$; with such a notion of stakes, the preferences $g, h \succ f \not\asymp g_\alpha h$ are compatible with representation (VII.1).

25. Schmeidler (1971) has shown that for incomplete transitive preferences (over appropriate spaces), the weak and strict preference orderings cannot both be continuous. Although his result does not apply here, due to the weakening of transitivity, it emphasises the subtlety of the issue of the continuity of the derived strict preference ordering. Note that this issue could also have been resolved by taking the strict preference ordering as primitive and using a version of the representation proposed by Bewley (2002); see Section VII.1.3.

26. 1_s is the characterisic function for s: $1_s(s') = 1$ for $s' = s$ and $1_s(s') = 0$ for $s' \neq s$. As specified in Section VII.2.1, c_s is the constant act taking the degenerate lottery yielding c_s for sure in all states.

Indifference consistency implies that indifferences cannot become strict preferences simply on altering the stakes. Under representation (VII.1), it is possible that $f \sim g$, because the expected utilities are equal for all $p \in D((f,g))$, but that $f_\alpha h \succ g_\alpha h$ for some α and h, because $D((f_\alpha h, g_\alpha h)) \supsetneq D((f,g))$ contains a probability measure q such that the expected utility of f calculated with q is strictly greater than that of g. In such cases, $D((f_\alpha h, g_\alpha h))$ is of a higher dimensionality than $D((f,g))$; the latter set, but not the former, is contained in the hyperplane defined by $u \circ f \sim u \circ g$. Indifference consistency rules out such possibilities; technically, it can be thought of as a constant dimensionality assumption on the sets of probability measures in the space $\Delta(\Sigma)$.

Full support is the behavioural formulation of the following full support property of Ξ: for each $s \in S$, there exists $b_s > 0$ such that $p(s) \geqslant b_s$ for all $p \in \bigcup_{\mathcal{C}_i \in \Xi} \mathcal{C}_i$. This property can be thought of as the analogue of full support for a probability measure, but for sets of measures and confidence rankings. In particular, it is stronger than simply asking that all probability measures in the confidence ranking have full support: it requires moreover that probability measures have a common non-zero lower bound on the values for each state.

VII.C Proofs

Throughout the Appendix, B will denote the space of all real-valued functions on S, and $ba(S)$ will denote the set of additive real-valued set functions on S, both under the Euclidean topology. B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $x \in \Re$, we define x^* to be the constant function taking value x.

VII.C.1 Proof of Theorem VII.1

The main part of the result is to show the sufficiency of the axioms for the representation (direction (i) to (ii)), the proof of which proceeds as follows. By standard arguments, we obtain a von Neumann-Morgernstern utility function on the consequences, which allows us to work with real-valued functions on S instead of acts. For each stakes level r, we define a preference relation $\preceq_r$ on these functions, which can be thought of as representing the preferences between corresponding acts considered "as if" the choices had stakes r. We show (Lemma VII.C.5) that, for each non-extremal stakes level r, $\preceq_r$ is a non-trivial, monotonic, affine, Archimedean pre-order, whence, by Gilboa et al. (2010, Corollary 1) (which is a version of a Ghirardato et al. (2004, Proposition A.2)), there is a closed convex set of probability measures $\mathcal{C}_r$ representing $\preceq_r$ according to the unanimity rule. Lemma VII.C.10 shows that the preference relations for minimal stakes can be represented according to the unanimity rule with the intersection of the $\mathcal{C}_r$ for the other stakes levels. By Lemma VII.C.7, the $\mathcal{C}_r$ form a nested family of sets, and we thus have a confidence ranking. By Lemmas VII.C.8 and VII.C.9, this confidence ranking is continuous. By construction, the function that assigns to any stakes level r the set $\mathcal{C}_r$ is a well-defined cautiousness coefficient.

Now we proceed with the proof. First we assume (i); we will show (ii). If $\leqq$ is trivial ($(f,g) \equiv (f',g')$ for all $f, f', g, g' \in \mathcal{A}$), then VII-A3 is equivalent to the standard transitivity axiom, and VII-A4 and VII-A6 are jointly equivalent to the standard independence axiom, so the result follows immediately from Gilboa et al. (2010, Theorem 1). We henceforth assume that $\leqq$ is not trivial. We begin with the following lemma.

Lemma VII.C.1. *There exists a non-constant utility function u representing the restriction of $\preceq$ to the constant acts.*

Proof. By VII-A1, VII-A2, VII-A4 and VII-A7, the restriction of $\preceq$ to constant acts is non-trivial complete, reflexive and satisfies independence and continuity. We now show that it is transitive. For any $c, c', c'' \in \Delta(X)$, suppose that $c \preceq c'$ and $c' \preceq c''$. If (c, c'') is $\leqq$-minimal, then VII-A3 immediately implies that $c \preceq c''$. Now suppose that (c, c'') is not $\leqq$-minimal. Consider firstly the case where $c \nsim c'$ and $c' \nsim c''$. By richness of $\leqq$, there exist $\alpha, \alpha' \in (0,1]$ and $h, h' \in \mathcal{A}$ such that $(c, c'') \leqq (c_\alpha h, c'_\alpha h)$ and $(c, c'') \leqslant (c'_{\alpha'} h', c''_{\alpha'} h')$. By VII-A1 and VII-A4, $c_\alpha h(s) \preceq c'_\alpha h(s)$ and $c'_{\alpha'} h'(s) \preceq c''_{\alpha'} h'(s)$ for all $s \in S$, from which it follows by VII-A5 that $c_\alpha h \preceq c'_\alpha h$ and $c'_{\alpha'} h' \preceq c''_{\alpha'} h'$. Hence, by VII-A3, $c \preceq c''$, as required. If $c \sim c'$ and $c' \nsim c''$, then by richness of $\leqq$, there exists $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $(c, c'') \leqq (c_\alpha h, c'_\alpha h)$; by VII-A1, VII-A4, VII-A5, and VII-A3 as previously, $c \preceq c''$. The case where $c' \sim c''$ and $c \nsim c'$ is treated similarly. The case where $c \sim c' \sim c''$ follows directly from VII-A3. The existence of u follows from the von Neumann-Morgenstern theorem. □

Let $K = u(\Delta(X))$ and $B(K)$ be the set of functions in B taking values in K. There is thus a many-to-one mapping between acts in $\mathcal{A}$ and elements of $B(K)$, given by $a = u \circ f$, for $f \in \mathcal{A}$. With slight abuse of notation, we use $\preceq$ to denote the order generated on $B(K)$

by $\preceq$ under this mapping, and $\leqq$ to denote the order generated on $B(K)\times B(K)$ by $\leqq$. ($\preceq$ and $\leqq$ are well-defined on $B(K)$ by VII-A5 and the extensionality of $\leqq$ respectively.)

Lemma VII.C.2. *For every $a, a', b, b' \in B(K)$ with $a \neq b$ and (a', b') not $\leqq$-minimal, there exists $\alpha \in (0, 1]$ and $l \in B(K)$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \equiv (a', b')$.*

Proof. If $(a, b) \equiv (a', b')$, then there is nothing to show. Suppose without loss of generality that $(a, b) < (a', b')$; the other case is treated similarly. By richness of $\leqq$, there exist $\beta \in (0, 1]$ and $l \in \mathcal{A}$ such that $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \geqq (a', b')$. If $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \equiv (a', b')$, then the result has been established; if not, then by continuity of $\leqq$, there exists $\alpha \in (\beta, 1)$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \equiv (a', b')$, as required. □

Let $\mathcal{S}$ be the set of equivalence classes of $\leqq$. As standard, $\leqq$ on $B(K)\times B(K)$ generates a relation on $\mathcal{S}$, which will be denoted $\leqslant$ (with symmetric and asymmetric components $=$ and $<$ respectively): for $r, s \in \mathcal{S}$, $r \leqslant s$ iff, for any $(f, g) \in r$ and $(f', g') \in s$, $(f, g) \leqq (f', g')$. $r \in \mathcal{S}$ is a minimal (respectively maximal) element if $r \leqslant s$ (resp. $r \geqslant s$) for all $s \in \mathcal{S}$. Note that, since $\leqslant$ is a linear ordering, there is at most one minimal (resp. maximal) element; if it exists, we denote the minimal (resp. maximal) element by $\underline{\mathcal{S}}$ (resp. $\overline{\mathcal{S}}$). $r \in \mathcal{S}$ is full if, for every $a, b \in B(K)$ with $a \neq b$, there exists $\alpha \in (0, 1]$ and $l \in B(K)$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$. It follows from Lemma VII.C.2 that every non-minimal element in $\mathcal{S}$ is full. Let $\mathcal{S}^f$ be the set of full elements in $\mathcal{S}$, and let $\mathcal{S}^+$ be the set of non-minimal full elements. For each $r \in \mathcal{S}^f$, let $\preceq_r$ be the reflexive binary relation on $B(K)$ such that, for all $a, b \in B(K)$ with $a \neq b$, $a \preceq_r b$ iff there exists $l \in B(K)$ and $\alpha \in (0, 1]$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$ and $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$. The following lemma implies that $a \preceq_r b$ iff $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ for every $l \in B(K)$ and $\alpha \in (0, 1]$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$.

Lemma VII.C.3. *For every $a, b, l, m \in B(K)$ and $\alpha, \beta \in (0, 1]$ with $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l), (\beta a + (1-\beta)m, \beta b + (1-\beta)m) \in r \in \mathcal{S}^f$, $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$.*

Proof. The result is an immediate consequence of VII-A5 if $a = b$, so consider $a, b \in B(K)$ with $a \neq b$. Without loss of generality, suppose that $\beta \leqslant \alpha$. Consider first the case where $\beta < \alpha$. Note that $\beta a + (1-\beta)m = \frac{\beta}{\alpha}(\alpha a + (1-\alpha)l) + (1 - \frac{\beta}{\alpha})(\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m)$, where

$\frac{\alpha\beta-\beta}{\alpha-\beta}l+\frac{\alpha-\alpha\beta}{\alpha-\beta}m \in B(K)$, since it is a $\frac{\alpha\beta-\beta}{\alpha-\beta}$-mix of l and m; similarly for $\beta b+(1-\beta)m$. Let $f, g, h \in \mathcal{A}$ be such that $\alpha a+(1-\alpha)l = u \circ f$, $\alpha b+(1-\alpha)l = u \circ g$ and $\frac{\alpha\beta-\beta}{\alpha-\beta}l+\frac{\alpha-\alpha\beta}{\alpha-\beta}m = u \circ h$; so $\beta a+(1-\beta)m = u \circ f_{\frac{\beta}{\alpha}}h$ and $\beta b+(1-\beta)m = u \circ g_{\frac{\beta}{\alpha}}c$. Since $(f, g) \equiv (f_{\frac{\beta}{\alpha}}h, g_{\frac{\beta}{\alpha}}h)$, by VII-A6, $f \asymp g$ iff $f_{\frac{\beta}{\alpha}}h \asymp g_{\frac{\beta}{\alpha}}h$. Hence, by VII-A4, $f \preceq g$ iff $f_{\frac{\beta}{\alpha}}h \preceq g_{\frac{\beta}{\alpha}}h$. So $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, as required.

Now consider the case where $\beta = \alpha$. If $l = m$, the result is immediate, so suppose that $l \neq m$. Note that if there exists $\epsilon \in (0, 1]$ and $k \in B(K)$ with $\epsilon \neq \alpha$ and $(\epsilon a + (1-\epsilon)k, \epsilon b + (1-\epsilon)k) \equiv (\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l)$, then, by applying the reasoning in the case above, we get $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $\epsilon a + (1-\epsilon)k \preceq \epsilon b + (1-\epsilon)k$ iff $\beta a+(1-\beta)m \preceq \beta b+(1-\beta)m$, as required. By the continuity and non-triviality of $\leqq$, $\mathcal{S}^f$ is not a singleton, so there exists $n \in B(K)$ and $\gamma \in (0, 1)$ such that $(\gamma(\alpha a+(1-\alpha)l)+(1-\gamma)n, \gamma(\alpha b+(1-\alpha)l)+(1-\gamma)n) \not\equiv (\alpha a+(1-\alpha)l, \alpha b+(1-\alpha)l)$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0, 1)$ and $n' \in B(K)$ such that $(\delta(\gamma(\alpha a + (1-\alpha)l) + (1-\gamma)n) + (1-\delta)n', \delta(\gamma(\alpha b + (1-\alpha)l) + (1-\gamma)n) + (1-\delta)n') \equiv (\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$. However, $\delta(\gamma(\alpha a + (1-\alpha)l) + (1-\gamma)n) + (1-\delta)n' = \alpha\gamma\delta a + (1-\alpha\gamma\delta)(\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l + \frac{1-\gamma}{1-\alpha\gamma}n) + \frac{1-\delta}{1-\alpha\gamma\delta}n')$, with $\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l + \frac{1-\gamma}{1-\alpha\gamma}n) + \frac{1-\delta}{1-\alpha\gamma\delta}n' \in B(K)$ since it is a mix of elements of $B(K)$. So $\alpha\gamma\delta$ and $\frac{\delta-\alpha\gamma\delta}{1-\alpha\gamma\delta}(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}l + \frac{1-\gamma}{1-\alpha\gamma}n) + \frac{1-\delta}{1-\alpha\gamma\delta}n'$ have the properties required above, and the result is established.

□

We now establish some properties of the relations $\preceq_r$.

Lemma VII.C.4. *For all $r, s \in \mathcal{S}^f$ with $r \geqslant s$, $\preceq_r \subseteq \preceq_s$.*

Proof. If $s = r$, there is nothing to show, so suppose not and consider $a, b \in B$ such that $a \preceq_r b$. If $a = b$, the result follows from the reflexivity of $\preceq_s$ and $\preceq_r$; henceforth suppose that this is not the case. Without loss of generality, it can be assumed that $(a, b) \in r$. (If not, replace a, b with $\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l$ where $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$ and continue as below.) It follows from Lemma VII.C.3 that $a \preceq b$. Let $\beta \in (0, 1)$ and $m \in B(K)$, be such that $(\beta a + (1-\beta)m, \beta b + (1-\beta)m) \in s$ (such β and m exist since $s \in \mathcal{S}^f$). VII-A6 and VII-A4 imply that $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, and hence $a \preceq_s b$, as required. □

Recall that a binary relation $\preceq$ on $B(K)$ is

— *non-trivial* if there exists $a, b \in B(K)$ such that $a \preceq b$ but not $a \succeq b$.

— *monotonic* if, for all $a, b, \in B(K)$, if $a \leqslant b$ then $a \preceq b$.
— *affine* if, for all $a, b, c \in B(K)$ and $\alpha \in (0,1)$, $a \preceq b$ iff $\alpha a + (1-\alpha)c \preceq \alpha b + (1-\alpha)c$.
— *Archimedean* if, for all $a, b, c \in B(K)$, the sets $\{\alpha \in [0,1] \mid \alpha a + (1-\alpha)b \succeq c\}$ and $\{\alpha \in [0,1] \mid \alpha a + (1-\alpha)b \preceq c\}$ are closed in $[0,1]$.
— a *pre-order* if $\preceq$ is reflexive and transitive.

Lemma VII.C.5. *For every $r \in \mathcal{S}^f$, $\preceq_r$ is a non-trivial, monotonic, affine pre-order. Moreover, if $r \in \mathcal{S}^+$, $\preceq_r$ is Archimedean.*

Proof. Non-triviality. By VII-A2, $\preceq$ is non-trivial; by VII-A5 and VII-A1, it follows that the restriction of $\preceq$ to $\Delta(X)$ is non-trivial. But $\preceq_r$ coincides with $\preceq$ on $\Delta(X)$, so it is non-trivial.

Monotonicity. Suppose that $a \leqslant b$ and $a \neq b$ (the result is immediate for $a = b$). Then, $\alpha a + (1-\alpha)l \leqslant \alpha b + (1-\alpha)l$ for $l \in B(K)$ and $\alpha \in (0,1]$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$. By monotonicity (VII-A5), $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$, and so $a \preceq_r b$.

Affineness. The result is immediate if $a = b$; henceforth suppose not. Since $r \in \mathcal{S}^f$, there exists $\beta \in (0,1]$ and $l \in B(K)$ such that $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \in r$. Consider $\beta(\alpha a + (1-\alpha)c) + (1-\beta)l$ and $\beta(\alpha b + (1-\alpha)c) + (1-\beta)l$: since $r \in \mathcal{S}^f$, there exists $\gamma \in (0,1]$ and $m \in B(K)$ such that $(\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m, \gamma(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m) \in r$. Note that $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m = \alpha\gamma(\beta a + (1-\beta)l) + (1-\alpha\gamma)(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m)$, where $\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m \in B(K)$ since it is a mix of elements of $B(K)$, and similarly for b. Let $f, g, h \in \mathcal{A}$ be such that $\beta a + (1-\beta)l = u \circ f$, $\beta b + (1-\beta)l = u \circ g$ and $\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m = u \circ h$. Since $(f, g) \equiv (f_{\alpha\gamma}h, g_{\alpha\gamma}h)$, by VII-A6 and VII-A4, $\beta a + (1-\beta)l \preceq \beta b + (1-\beta)l$ iff $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m \preceq \gamma(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m$. But since $\gamma(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\gamma)m = \beta\gamma(\alpha a + (1-\alpha)c) + (1-\beta\gamma)(\frac{\gamma-\beta\gamma}{1-\beta\gamma}l + \frac{1-\gamma}{1-\beta\gamma}m)$, and similarly for b, it follows that $a \preceq_r b$ iff $\alpha a + (1-\alpha)c \preceq_r \alpha b + (1-\alpha)c$, as required.

Pre-order. Reflexivity follows from the definition of $\preceq_r$. As for transitivity, suppose that $a \preceq_r b$ and $b \preceq_r c$ and that $a \neq b \neq c$ (if $a = b$, $b = c$ or $a = c$, the result is immediate). Since $r \in \mathcal{S}^f$, there exists $l \in B(K)$ and $\alpha \in (0,1]$ such that $(\alpha a + (1-\alpha)l, \alpha c + (1-\alpha)l) \in r$. Moreover, there exists $m, n \in B(K)$ and $\beta, \gamma \in (0,1]$ such that $(\beta(\alpha a + (1-\alpha)l) + (1-\beta)m), \beta(\alpha b + (1-\alpha)l) + (1-\beta)m)) \in r$ and $(\gamma(\alpha b + (1-\alpha)l) + (1-\gamma)m), \gamma(\alpha c + (1-\alpha)l) + (1-\gamma)m) \in r$. Since $a \preceq_r b$ and $b \preceq_r c$,

$\beta(\alpha a+(1-\alpha)l)+(1-\beta)m \preceq \beta(\alpha b+(1-\alpha)l)+(1-\beta)m)$ and $\gamma(\alpha b+(1-\alpha)l)+(1-\gamma)m \preceq \gamma(\alpha c+(1-\alpha)l)+(1-\gamma)m$. Hence, by VII-A3, $\alpha a+(1-\alpha)l \preceq \alpha c+(1-\alpha)l$, and so $a \preceq_r c$, as required.

Archimedean. Let r be a non-minimal element of $\mathcal{S}^f$. Consider $\{\alpha \in [0,1] \mid \alpha a+(1-\alpha)b \succeq_r c\}$; the other case is dealt with similarly. Let $\bar{\alpha}$ be a limit point of this set, and without loss of generality, assume that $(\bar{\alpha}a+(1-\bar{\alpha})b, c) \in r$ (if not, replace a, b, c with appropriate mixtures for which this is the case). It needs to be shown that $\bar{\alpha}a+(1-\bar{\alpha})b \succeq_r c$. If $\bar{\alpha}a+(1-\bar{\alpha})b = c$ the result is immediate; suppose henceforth that this is not the case. If there is a open interval I in $\{\alpha \in [0,1] \mid \alpha a+(1-\alpha)b \succeq_r c\}$ such that $\bar{\alpha}$ is a limit point of I and such that $(\beta a+(1-\beta)b, c) \leqq (\bar{\alpha}a+(1-\bar{\alpha})b, c)$ for all $\beta \in I$, then, by Lemma VII.C.4, $\beta a+(1-\beta)b \succeq c$ for all $\beta \in I$, whence $\bar{\alpha}a+(1-\bar{\alpha})b \succeq c$ by VII-A7, and so $\bar{\alpha}a+(1-\bar{\alpha})b \succeq_r c$ as required.

Now suppose that there is no such interval. Since r is a non-minimal element of $\mathcal{S}$, by the Continuity of $\leqq$ and Lemma VII.C.2, there exists $l \in B(K)$ and $\bar{\delta} \in (0,1)$ such that $(\bar{\delta}(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\bar{\delta})l, \bar{\delta}c+(1-\bar{\delta})l) < (\bar{\alpha}a+(1-\bar{\alpha})b, c)$. Let $\gamma = \min\{\delta \in (\bar{\delta}, 1] \mid (\delta(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\delta)l, \delta c+(1-\delta)l) \geqq (\bar{\alpha}a+(1-\bar{\alpha})b, c)\}$ (by Continuity of $\leqq$ this is a minimum). Consider any $\delta \in (\bar{\delta}, \gamma)$; by the definition of γ, $(\delta(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\delta)l, \delta c+(1-\delta)l) < (\bar{\alpha}a+(1-\bar{\alpha})b, c)$. Note moreover that $\delta(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\delta)l = \bar{\alpha}(\delta a+(1-\delta)l)+(1-\bar{\alpha})(\delta b+(1-\delta)l)$. So, by the Continuity of $\leqq$, there is an open interval $I_\delta \subseteq (0,1)$ containing $\bar{\alpha}$ such that, for all $\beta \in I_\delta$, $(\beta(\delta a+(1-\delta)l)+(1-\beta)(\delta b+(1-\delta)l), \delta c+(1-\delta)l) < (\bar{\alpha}a+(1-\bar{\alpha})b, c)$. Note that $I_\delta \cap \{\alpha \in [0,1] \mid \alpha a+(1-\alpha)b \succeq_r c\}$ is non-empty, since $\bar{\alpha}$ is a limit point of $\{\alpha \in [0,1] \mid \alpha a+(1-\alpha)b \succeq_r c\}$. Furthermore, since $\beta(\delta a+(1-\delta)l)+(1-\beta)(\delta b+(1-\delta)l) = \delta(\beta a+(1-\beta)b)+(1-\delta)l$, Lemma VII.C.4 implies that $\beta(\delta a+(1-\delta)l)+(1-\beta)(\delta b+(1-\delta)l) \succeq \delta c+(1-\delta)l$ for all $\beta \in I_\delta \cap \{\alpha \in [0,1] \mid \alpha a+(1-\alpha)b \succeq_r c\}$. It follows by VII-A7 that $\bar{\alpha}(\delta a+(1-\delta)l)+(1-\bar{\alpha})(\delta b+(1-\delta)l) \succeq \delta c+(1-\delta)l$. Since this holds for all $\delta \in (\bar{\delta}, \gamma)$, it follows by VII-A7 that $\gamma(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\gamma)l \succeq \gamma c+(1-\gamma)l$; whence, since $(\gamma(\bar{\alpha}a+(1-\bar{\alpha})b)+(1-\gamma)l, \gamma c+(1-\gamma)l) \in r$, $\bar{\alpha}a+(1-\bar{\alpha})b \succeq_r c$, as required.

□

Lemma VII.C.6. *For each $r \in \mathcal{S}^+$, there exists a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $a, b \in B$, $a \preceq_r b$ iff*

(VII.C.1) $$\sum_{s \in S} a(s)p(s) \leqslant \sum_{s \in S} b(s)p(s) \quad \textit{for all } p \in \mathcal{C}_r$$

Proof. This follows from Lemma VII.C.5, by Gilboa et al. (2010, Corollary 1),[27] which establishes such a representation for non-trivial, monotonic, affine, Archimedean pre-orders. □

Lemma VII.C.7. *For all $r, s \in \mathcal{S}^+$ with $r \geqslant s$, $\mathcal{C}_s \subseteq \mathcal{C}_r$.*

Proof. This follows directly from Lemma VII.C.4 and Ghirardato et al. (2004, Proposition A.1). □

Lemma VII.C.8. *For all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$.*

Proof. By Lemma VII.C.7, $\mathcal{C}_r \supseteq \mathcal{C}_{r'}$ for all $r' < r$. Suppose, for reductio, that $\mathcal{C}_r \supsetneq \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$, so that there exists a point (probability measure) $p \in \mathcal{C}_r \backslash \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0, 1]$ and $m \in B(K)$ such that $(\delta a + (1 - \delta)m, \delta\alpha^* + (1 - \delta)m) \in r$. By Lemma VII.C.2, there exists $l \in B(K)$ and $\beta \in (0, 1)$ such that $(\beta(\delta a + (1 - \delta)m) + (1 - \beta)l, \beta(\delta\alpha^* + (1 - \delta)m) + (1 - \beta)l) < (\delta a + (1 - \delta)m, \delta\alpha^* + (1 - \delta)m)$. Let $\beta' = \min\{\gamma \in [\beta, 1] | \ (\beta(\delta a + (1 - \delta)m) + (1 - \beta)l, \beta(\delta\alpha^* + (1 - \delta)m) + (1 - \beta)l) \geqq (\delta a + (1 - \delta)m, \delta\alpha^* + (1 - \delta)m)\}$ (this is a minimum by the Continuity of $\leqq$). Taking $f, g, h \in \mathcal{A}$ such that $u \circ f = \delta a + (1 - \delta)m$, $u \circ g = \delta\alpha^* + (1 - \delta)m$ and $u \circ h = l$, it follows, by the construction, that for any $\gamma \in (\beta, \beta')$, $g_\gamma h \preceq f_\gamma h$. However, by construction, $g_{\beta'} h \npreceq f_{\beta'} h$, contradicting VII-A7. Hence $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. □

Lemma VII.C.9. *For all non-maximal $r \in \mathcal{S}^+$, $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.*

Proof. By Lemma VII.C.7, $\mathcal{C}_r \subseteq \mathcal{C}_{r'}$ for all $r' > r$. Suppose, for reductio, that $\mathcal{C}_r \subsetneq \bigcap_{r'>r} \mathcal{C}_{r'}$, so that there exists a point (probability measure) $p \in \bigcap_{r'>r} \mathcal{C}_{r'} \backslash \mathcal{C}_r$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$, an $\alpha \in \Re$ and an $\epsilon > 0$

27. See Ghirardato et al. (2004, Proposition A.2) for a related result.

such that $\phi(p) \leqslant \alpha - \epsilon$ and $\alpha \leqslant \phi(q)$ for all $q \in \mathcal{C}_r$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1]$ and $m \in B(K)$ such that $(\delta a + (1-\delta)m, \delta\alpha^* + (1-\delta)m) \in r$. Take any $x \in K$ with $x \leqslant \alpha, a(s)$ for all $s \in S$, and let $f, g, h \in \mathcal{A}$ be such that $u \circ f = \delta a + (1-\delta)m$, $u \circ g = \delta\alpha^* + (1-\delta)m$, $u \circ h = \delta x^* + (1-\delta)m$. Let $\beta \in (0,1)$ be such that $u \circ g_\beta h = \delta(\alpha - \frac{\epsilon}{2})^* + (1-\delta)m$; such a β exists by the definition of g and h. By construction, $f \succeq g$, $f(s), g(s) \succeq h(s)$ for all $s \in S$, there exists $\alpha \in (0,1]$ and $e \in \mathcal{A}$ with $(f_\alpha e, (g_\beta h)_\alpha e) > (f,g)$ and $(f_\alpha e, (g_\beta h)_\alpha e)$ not $\leqq$-maximal, but, for all such α and e, $f_\alpha e \not\succeq (g_\beta h)_\alpha e$. Hence, by VII-A6, for all $\alpha \in (0,1]$ and $e \in \mathcal{A}$ such that $(f_\alpha e, (g_\beta h)_\alpha e) > (f,g)$, $f_\alpha e \not\succeq (g_\beta h)_\alpha e$, contradicting VII-A8; hence $\mathcal{C}_r = \bigcap_{r' > r} \mathcal{C}_{r'}$.

□

Lemma VII.C.10. *Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on $B(K)$ generated by (VII.C.1) with the set of probability measures $\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r$. If there exists a non-empty minimal element of $\mathcal{S}$, $\underline{\mathcal{S}}$, then $\preceq|_{\underline{\mathcal{S}}} = \preceq_{\bigcap \mathcal{S}}|_{\underline{\mathcal{S}}}$.*

Proof. Let $a, b \in B(K)$ be such that $(a,b) \in \underline{\mathcal{S}}$ and suppose that $a \preceq_{\bigcap \mathcal{S}} b$. Let $x = \min\{a(s), b(s) \mid s \in S\}$. Since $a \preceq_{\bigcap \mathcal{S}} b$, it follows from representation (VII.C.1) and Lemma VII.C.6 that for each $\beta \in (0,1)$, there exists a non-maximal $s > \underline{\mathcal{S}}$ such that $\beta a + (1-\beta)x^* \preceq_s b$, and thus, by Lemma VII.C.2, there exists $\alpha \in (0,1]$ and $e \in \mathcal{A}$ such that $\alpha(\beta a + (1-\beta)x^*) + (1-\alpha)e \preceq \alpha b + (1-\alpha)e$. Hence, by VII-A8, $a \preceq b$, as required. Now suppose that $a \preceq b$. By VII-A8, for every $\beta \in (0,1)$, there exists $r > \underline{\mathcal{S}}$ such that $b \succeq_r \beta a + (1-\beta)x^*$, where x is as defined above. So, by Lemma VII.C.6, $b \succeq_{\bigcap \mathcal{S}} \beta a + (1-\beta)x^*$ for all $\beta \in (0,1)$. Since $\preceq_{\bigcap \mathcal{S}}$ is Archimedean, it follows that $b \succeq_{\bigcap \mathcal{S}} a$, as required.

□

Conclusion of the proof of Theorem VII.1. Define

$$\Xi = \begin{cases} \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} & \text{if } \mathcal{S} = \mathcal{S}^+ \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r\} & \text{if } \mathcal{S} = \mathcal{S}^+ \cup \{\underline{\mathcal{S}}\} \end{cases}$$

where the $\mathcal{C}_r$ are as specified in Lemma VII.C.6. It follows from Lemma VII.C.7 that Ξ is a nested family of sets. Since the $\mathcal{C}_r$ are closed and convex for all $r \in \mathcal{S}^+$ (Lemma

VII.C.6), Ξ is a confidence ranking. By Lemmas VII.C.8 and VII.C.9, Ξ is continuous. D is defined as follows: for all $(f, g) \in \mathcal{A} \times \mathcal{A}$, if $[(f, g)] \in \mathcal{S}^+$, then $D((f, g)) = \mathcal{C}_{[(f,g)]}$, and if $(f, g) \in \underline{\mathcal{S}}$, then $D((f, g)) = \bigcap_{s \in \mathcal{S}^+} \mathcal{C}_s$. Order preservation and surjectivity of D are immediate from the definition and Lemma VII.C.7. By construction and Lemma VII.C.10, u, Ξ, D represent $\preceq$ according to (VII.1).

Now consider the clause concerning the centering axiom VII-A9. Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on $B(K)$ generated by (VII.C.1) with the set of probability measures $\bigcap_{r \in \mathcal{S}^f} \mathcal{C}_r$. Obviously $\preceq_{\bigcap \mathcal{S}} \supseteq \bigcup_{r \in \mathcal{S}^f} \preceq_r$. By VII-A9, for each $f \in \mathcal{A}$, there is at most one $c \in \Delta(X)$ for which (f, c) is not ordered by $\bigcup_{r \in \mathcal{S}^f} \preceq_r$; since $\preceq_{\bigcap \mathcal{S}}$ is Archimedean, it follows that (f, c) is ordered by $\preceq_{\bigcap \mathcal{S}}$ for every $f \in \mathcal{A}$ and $c \in \Delta(X)$, and hence that $\preceq_{\bigcap \mathcal{S}}$ is complete. It follows from the form of (VII.C.1) that $\bigcap_{r \in \mathcal{S}^f} \mathcal{C}_r$ is a singleton, as required.

The direction from (ii) to (i) is generally straightforward. The only interesting case is continuity (VII-A7). Consider any $f, g, h \in \mathcal{A}$, and the set $\{(\alpha, \beta) \in [0, 1]^2 | \ f_\alpha h \preceq g_\beta h\}$. Suppose that (α^*, β^*) is a limit point of this set, and consider a sequence $((\alpha_i, \beta_i))$ of members of the set with $(\alpha_i, \beta_i) \to (\alpha^*, \beta^*)$. If there exists a subsequence of $((\alpha_i, \beta_i))$, tending to (α^*, β^*), such that $(f_{\alpha_{i_n}} h, g_{\beta_{i_n}} h) \geqq (f_{\alpha^*} h, g_{\beta^*} h)$ for all $(\alpha_{i_n}, \beta_{i_n})$ in this sequence, then the case can be treated in a similar way to that used below, relying on the continuity of the confidence ranking. Henceforth, we consider the case where there is no such sequence. In this case, there exists $f', g' \in \mathcal{A}$ with $(f', g') < (f_{\alpha^*} h, g_{\beta^*} h)$. Moreover, by the continuity of $\leqq$, for each such (f', g'), there is an open interval around (α^*, β^*) such that $(f_\gamma h, g_\delta h) \geqq (f', g')$ for any (γ, δ) in this interval. Hence, for each such (f', g'), there is a subsequence $((\alpha_{j_n^{[(f',g')]}}, \beta_{j_n^{[(f',g')]}}))$ of $((\alpha_i, \beta_i))$, tending to (α^*, β^*), with $(f_{\alpha_{j_n^{[(f',g')]}}} h, g_{\beta_{j_n^{[(f',g')]}}} h) \geqq (f', g')$ for all $n \in \mathbb{N}$. It follows, since D is order-preserving, that for all $n \in \mathbb{N}$, $\sum_{s \in S} u(f_{\alpha_{j_n^{[(f',g')]}}} h(s)).p(s) \leqslant \sum_{s \in S} u(g_{\beta_{j_n^{[(f',g')]}}} h(s)).p(s)$, for all $p \in D((f', g'))$. Hence, by the continuity of the representation, it follows that $\sum_{s \in S} u(f_{\alpha^*} h(s)).p(s) \leqslant \sum_{s \in S} u(g_{\beta^*} h(s)).p(s))$, for all $p \in D((f', g'))$. Since this holds for every $(f', g') < (f, g)$, and since, by the continuity of the confidence ranking and the surjectivity of D, $D((f, g)) = \overline{\bigcup_{(f',g')<(f,g)} D((f', g'))}$, there cannot be a $q \in D((f, g))$ such that $\sum_{s \in S} u(f_{\alpha^*} h(s)).q(s) > \sum_{s \in S} u(g_{\beta^*} h(s)).q(s)$. So $f_{\alpha^*} h \preceq g_{\beta^*} h$, and hence $\{(\alpha, \beta) \in [0, 1]^2 | \ f_\alpha h \preceq g_\beta h\}$ is closed, as required.

Finally, consider the uniqueness clause. Uniqueness of u follows from the von Neumann-Morgenstern theorem. As regards uniqueness of Ξ, proceed by reductio; suppose that

u, Ξ_1, D_1 and u, Ξ_2, D_2 both represent $\preceq$ according to (VII.1), with $\Xi_1 \neq \Xi_2$. Since Ξ_1 and Ξ_2 are continuous, they must differ on some non-minimal element; hence, by the surjectivity of the D_i, there exists $(f, g) \in \mathcal{A} \times \mathcal{A}$ with non-minimal stakes such that $D_1((f, g)) \neq D_2((f, g))$. Suppose, without loss of generality, that $p \in D_1((f, g)) \backslash D_2((f, g))$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in D_2((f, g))$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ, a and α can be chosen so that $\alpha \in K$, $a \in B(K)$ and $(a, \alpha^*) \equiv (u \circ f, u \circ g)$. Taking $h \in \mathcal{A}$ such that $u \circ h = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have that $\sum_{s \in S} u(h(s))p(s) \geqslant \sum_{s \in S} u(c)p(s)$ for all p s.t. $p \in D_2((h, c))$, whereas this is not the case for all p s.t. $p \in D_1((h, c))$, contradicting the assumption that both u, Ξ_1, D_1 and u, Ξ_2, D_2 represent $\preceq$. A similar argument establishes the uniqueness of D.

□

VII.C.2 Proofs of results in Sections VII.3.3 and VII.4

Proof of Proposition VII.1. Let the assumptions of the Proposition be satisfied. (ii) implies (i) is straightforward, so we consider only (i) implies (ii). Since the preference relations are complete on the set of constant acts, they coincide on that set; hence, by the uniqueness clause of the von Neumann-Morgenstern theorem, u_2 is a positive affine transformation of u_1. Hence K and the mapping from $\mathcal{A}$ to $B(K)$ used in the proof of Theorem VII.1 can be taken to be the same for the two agents; we use the notation employed in that proof. By (i), for every $r \in \mathcal{S}^+$, $\preceq_r^1 \subseteq \preceq_r^2$, and so, by Ghirardato et al. (2004, Proposition A.1), $\mathcal{C}_r^2 \subseteq \mathcal{C}_r^1$. It follows that $\Xi_2 \sqsubseteq \Xi_1$ and $D_2((f, g)) \subseteq D_1((f, g))$ for all $(f, g) \in \mathcal{A} \times \mathcal{A}$. □

Proof of Proposition VII.2. The 'if' direction is straightforward. The 'only if' direction is a simple corollary of the proof of Theorem VII.1. On the one hand, if $\preceq^1$ and $\preceq^2$ are confidence equivalent, they have identical preferences over constant acts (of which they are maximally confident), and hence the same utilities up to positive affine transformation. On the other hand, if they are confidence equivalent, the sets of preferences $\{\preceq_r \, | r \in \mathcal{S}^+\}$ defined in the proof of Theorem VII.1 are the same, and so the confidence rankings are the same. □

Proof of Theorem VII.2. First consider part (ii). Showing the necessity of the axioms is straightforward; we now show sufficiency. Take $f, g \in \mathcal{A}$, and consider firstly the case where $\min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s) \geqslant \min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s)$. Suppose there exist $c \in \Delta(X)$, $\alpha \in (0,1]$, $h \in \mathcal{A}$ such that $(f_\alpha h, c_\alpha h) \geqq (f, g)$ and $f_\alpha h \succeq c_\alpha h$; it thus follows that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(c)$ for all $p \in D((f,g))$. If (f,g) is not $\leqq$-minimal, then, by the richness of $\leqq$, for each $d \in \Delta(X)$ with $d \prec c$, there exist $\alpha' \in (0,1]$, $h' \in \mathcal{A}$ such that $(g_{\alpha'}h', d_{\alpha'}h') \equiv (f,g)$; moreover, for any such d, h', α', $g_{\alpha'}h' \succeq d_{\alpha'}h$. If (f,g) is $\leqq$-minimal, then, by the continuity of Ξ and the subjectivity of D, for each $d \in \Delta(X)$ with $d \prec c$, there exists $(f', g') > (f,g)$ such that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(d)$ for all $p \in D((f', g'))$. Moreover, by the richness of $\leqq$, there exist $\alpha' \in (0,1]$, $h' \in \mathcal{A}$ such that $(g_{\alpha'}h', d_{\alpha'}h') \equiv (f', g')$; for any such d, h', α', $g_{\alpha'}h' \succeq d_{\alpha'}h$. Hence, for any $d \in \Delta(X)$ with $d \prec c$, there exist $\alpha' \in (0,1]$, $h' \in \mathcal{A}$ such that $(g_{\alpha'}h', d_{\alpha'}h') \geqq (f,g)$ and $g_{\alpha'}h' \succeq d_{\alpha'}h$. It follows by VII-B2^S that $g \nprec^c f$, and thus, by VII-B1, $g \succeq^c f$.

Now suppose that $\min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s) < \min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s)$, so there exist $c, d \in \Delta(X)$ with $\min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s) < u(d) < u(c) < \min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s)$. Take any such c, d. First, note that it is not the case that $\sum_{s\in S} u(g(s)).p(s) \leqslant u(d)$ for all $p \in D((f,g))$. It thus follows from representation (VII.1) that for all $h' \in \mathcal{A}$ and $\alpha' \in [0,1]$ such that $(g_{\alpha'}h', d_{\alpha'}h') \geqq (f,g)$, $g_{\alpha'}h' \nsucceq^d d_{\alpha'}h'$. However, we have $u(c) \leqslant \sum_{s\in S} u(f(s)).p(s)$ for all $p \in D((f,g))$; moreover, if (f,g) is not $\leqq$-maximal, it follows from the continuity of Ξ and the subjectivity of D that there exists $(f', g') > (f,g)$ such that $u(c) \leqslant \sum_{s\in S} u(f(s)).p(s)$ for all $p \in D((f', g'))$. If (f,g) is not $\leqq$-minimal, then by the richness of $\leqq$, there exist $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $(f_\alpha h, c_\alpha h) \equiv (f,g)$; by representation (VII.1), $f_\alpha h \succeq c_\alpha h$. If (f,g) is $\leqq$-minimal, then take any $(f', g') > (f,g)$ such that $u(c) \leqslant \sum_{s\in S} u(f(s)).p(s)$ for all $p \in D((f', g'))$; by the richness of $\leqq$, there exist $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $(f_\alpha h, c_\alpha h) \equiv (f', g') > (f,g)$, and by representation (VII.1), $f_\alpha h \succeq c_\alpha h$. It follows from VII-B2^S that $g \prec^c f$, as required. Hence representation (VII.3) holds.

A similar argument establishes part (i). □

VII.C.3 Proofs of results in Section VII.5

As stated in Section VII.5 (see in particular footnote 15), we continue to use the standard notation (and generic terms $f, g \ldots$) for acts, as well as the standard notation (and generic terms $x, z \ldots$) for commodities. In particular, for commodities x, z and $\alpha \in [0,1]$,

$\alpha x+(1-\alpha)z$ is the standard vector sum of products of the two commodities, whereas $x_\alpha z$ is the act obtained by applying the mixture operation on (the acts corresponding to) the commodities. Whilst $x_\alpha z$ does not in general belong to $\Re^S_+$, for any preference relation $\preceq^i$ with the properties specified in Section VII.5, there is a natural element in $\Re^S_+$ corresponding to it; namely $((u^i)^{-1}(\alpha u^i(x_1) + (1-\alpha)u^i(z_1)), \ldots, (u^i)^{-1}(\alpha u^i(x_{|S|}) + (1-\alpha)u^i(z_{|S|})))$. (At each state, the lottery obtained in that state is replaced by its certainty equivalent.) Henceforth we denote this element by $x^i_\alpha z$.

We first require the following Lemma.

Lemma VII.C.11. *Under Assumption VII.1, the strict preferences $\succ^i$ have the following* reduced convexity *property: for all $f, g, h \in \mathcal{A}$, if $g, h \succ^i f$, then, for all $\alpha \in (0,1)$, there exists $\beta \in (0,1]$ such that $(g_\alpha h)_{\beta'} f \succ^i f$ for all $\beta' \in (0,\beta]$.*

Proof. Let $f, g, h \in \mathcal{A}$ such that $g, h \succ^i f$, and consider $\alpha \in (0,1)$. By the monotone decreasing and continuity properties of stakes, there exists $\beta \in (0,1]$ such that $((g_\alpha h)_\beta f, f) \leqq \min\{(g,f),(h,f)\}$. By the representation (VII.1), it follows that $(g_\alpha h)_\beta f \succeq^i f$. However, if $(g_\alpha h)_\beta f \sim^i f$, then, by indifference consistency (and the properties of the representation) $(g_\alpha h)_\gamma h' \sim^i f_\gamma h'$ for any $\gamma \in (0,1]$, $h' \in \mathcal{A}$ with $((g_\alpha h)_\gamma h', f_\gamma h') \equiv \min\{(g,f),(h,f)\}$, contradicting $g, h \succ^i f$, under representation (VII.1). Hence $(g_\alpha h)_\beta f \succ^i f$ and, by indifference consistency, the fact that the stakes are monotone decreasing, and the properties of the representation, $(g_\alpha h)_{\beta'} f \succ^i f$ for all $\beta' \in (0,\beta]$, as required. □

Proof of Theorem VII.3. For any $x \in \Re^S_+$, let $\pi^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+, \text{if } z \succ x, \text{ then } p \cdot z > p \cdot x\}$, and let $\bar{\pi}^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+, \text{if } z \succ x, \text{ then } p \cdot z \geqslant p \cdot x\}$. On inspection, it is straightforward to check that the reduced convexity property (Lemma VII.C.11), combined with the concavity of u and the monotonicity of representation (VII.1), is sufficient for the application of standard arguments on welfare theorems in the absence of completeness and transitivity, notably Fon and Otani (1979), yielding the conclusion that, if x is Pareto optimal, there exists $p \in \bigcap_i \bar{\pi}^i(f^i)$. (In a word, in the presence of reduced convexity and concavity of the utility function, Pareto optimality implies that the convex hull of the strict upper contour set of x^i is disjoint from $\{x^i\}$, allowing application of a separating hyperplane theorem. By monotonicity of representation (VII.1), the separating hyperplane has a positive normal; by normalising, this yields a $p \in \bigcap_i \bar{\pi}^i(x^i)$.) We show that $\pi^i(x) = \bar{\pi}^i(x)$ for all i and $x \in \Re^S_+$. Suppose not, and let $p \in \bar{\pi}^i(x) \backslash \pi^i(x)$

for some i and x; so there exists z with $z \succ^i x$ and $p \cdot z = p \cdot x$. By the fact that stakes are monotone decreasing, indifference consistency of preferences, and representation (VII.1), $z^i_\alpha x \succ^i x$ for any $\alpha \in (0, 1]$. By strict concavity of u, for all $s \in S$, $(z^i_\alpha x)_s = (u^i)^{-1}(\alpha u^i(z_s) + (1-\alpha)u^i(x_s)) \leqslant \alpha z_s + (1-\alpha)x_s$, with strict inequality whenever $z_s \neq x_s$. It follows that either $p \cdot x = p \cdot (\alpha z + (1-\alpha)x) > p \cdot (z^i_\alpha x)$, contradicting the assumption that $p \in \bar{\pi}^i(f)$, or $p(s) = 0$ whenever $x_s \neq z_s$. Consider the latter case, and let $S_1 = \{s \in S | \, p(s) = 0\}$. By full support, $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \backslash S_1)} > 0$; pick any $\delta > 0$ with $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \backslash S_1)} > \delta \, \frac{\max_{s \notin S_1} u^{i\prime}(z_s)}{\min_{s \in S_1} u^{i\prime}(z_s)}$. For $\epsilon > 0$ and define the allocation z^ϵ as follows: $z^\epsilon_s = \epsilon$ for $s \in S_1$, and $z^\epsilon_s = -\epsilon.\delta$ for $s \notin S_1$. By the definition of z^ϵ, $z + z^\epsilon \succ^i x$ for ϵ sufficiently small, and $p \cdot (z + z^\epsilon) < p \cdot x$ for all $\epsilon > 0$, contradicting the assumption that $p \in \bar{\pi}^i(x)$. Hence $\bar{\pi}^i(x) = \pi^i(x)$ as required.

By standard arguments, if $\bigcap_i \pi^i(x^i) \neq \emptyset$, then $(x^1, \ldots, x^n)$ is Pareto optimal. It remains to show that $\Pi^i(x) = \pi^i(x)$ for all i and $x \in \Re^S_+$.

We first show that $\Pi^i(x) \subseteq \pi^i(x)$. Note that, if $z \succ^i x$, then $\sum_s p(s)(u^i(z_s) - u^i(x_s)) > 0$ for all $p \in ri(D^i((x, z)))$ and hence for all $p \in \bigcap_{z \neq x} ri(D^i(x, z))$. By concavity of u^i, it follows that $\sum_s p(s)u^{i\prime}(x_s)(z_s - x_s) > 0$ for all $p \in \bigcap_{z \neq x} ri(D^i(x, z))$. Renormalising, it follows that, for any $q \in \Pi^i(f)$, $\sum_s q_s \cdot (z_s - x_s) > 0$, and hence that $q \cdot z > q \cdot x$. Since this holds for all $z \in \Re^S_+$ with $z \succ^i x$, $q \in \pi^i(x)$.

We now show that $\pi^i(x) \subseteq \Pi^i(x)$. Suppose not, and let $\bar{p} \in \pi^i(x) \backslash \Pi^i(x)$. Since, as is straightforwardly checked, $\Pi^i(x)$ is convex, by a separation theorem, there exists $y \in \Re^S$ and $b \in \Re$ with $\bar{p} \cdot y \leqslant b \leqslant q \cdot y$ for all $q \in \overline{\Pi^i(x)}$ where the right hand inequality is strict for all $q \in ri(\Pi^i(x))$. Without loss of generality, we can take $b = 0$. Since this implies that, for $\alpha > 0$, $q \cdot \alpha y = \frac{1}{\sum_{t \in S} p(t) u^{i\prime}(x_t)} \sum_s p(s) u^{i\prime}(x_s) \alpha y_s \geqslant 0$ for all $p \in \overline{\bigcap_{z \neq x} ri(D^i(x, z))}$ with strict inequality for all $p \in ri(\bigcap_{z \neq x} ri(D^i(x, z)))$, and since $\overline{\bigcap_{z \neq x} ri(D^i(x, z))}$ is compact, it follows that, for α sufficiently small, $\sum_s p(s)(u^i((x + \alpha y)_s) - u^i(x_s)) \geqslant 0$ for all $p \in \bigcap_{z \neq x} ri(D^i(x, z))$, with strict inequality for some such p. Let A be the set of α possessing this property and such that $x + \alpha y \in \Re^S_+$; we show that $x + \alpha y \succ^i x$ for some $\alpha \in A$. If not, then for every $\alpha \in A$, there exists $\hat{p} \in D^i((x + \alpha y, x))$ with $\sum_s \hat{p}(s)(u^i((x + \alpha y)_s) - u^i(x_s)) < 0$. By nestedness of Ξ^i, it follows that there exists $\hat{p} \in \bigcap_{z \neq x} ri(D^i(x, z))$ with $\sum_s \hat{p}(s)(u^i((x + \alpha y)_s) - u^i(x_s)) < 0$, contradicting the inverse inequality above. Hence $x + \alpha y \succ^i x$ for some $\alpha > 0$, whereas $\bar{p} \cdot (x + \alpha y) \leqslant \bar{p} \cdot x$, so $\bar{p} \notin \pi^i(x)$, as required.

□

Proof of Proposition VII.3. We show this on the example given in the text. Consider the full insurance allocation $(z^1_{\bar{\delta}}, z^2_{\bar{\delta}}) = ((\bar{\delta}w, \bar{\delta}w), ((1-\bar{\delta})w, (1-\bar{\delta})w)$. Agent 2 would accept to exchange x^2 for this ($x^2 \prec^2 z^2_{\bar{\delta}}$) iff:

$$0.5u^2((1-\delta)w) + 0.5u^2(\delta w) < u^2(1-\bar{\delta})$$

This gives a strict upper bound ν on $\bar{\delta}$. If a condition analogous to that in (VII.5) holds, namely:

$$\text{(VII.C.2)} \qquad \min\{\eta w \max\{|\delta - \bar{\delta}|, |\bar{\delta} - (1-\delta)|\}, 0.45\} + 0.5 > \frac{\bar{\delta}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}$$

for all $\bar{\delta} < \nu$, then, for all $\bar{\delta}$ such that $x^2 \prec^2 z^2_{\bar{\delta}}$, $x^1 \nprec^1 z^1_{\bar{\delta}}$; hence there are no Pareto optimal allocations accessible from (x^1, x^2). It is straightforwardly checked that, with the parameter values given in the text, these conditions are satisfied.

□

Proof of Proposition VII.4. It suffices to give an example where no Pareto optimum is m-accessible for any finite m; we use a refinement of the previous example, with $\gamma^1 = \gamma^2 = 1$. Take an allocation $(x^1, x^2) = ((\delta_1 w, \delta_2 w), ((1-\delta_1)w, (1-\delta_2)w))$ with the following properties:

(a) $1 > \delta_1 > \delta_2 > 0$

(b) $\delta_1 - \delta_2 < 18\delta_2(1-\delta_1)$

(c) $\eta w > \max\left\{\frac{1}{(1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1)}, \frac{1}{\delta_1+\delta_1^{0.5}\delta_2^{0.5}-2\delta_1^{0.5}\delta_2^{1.5}}, 2\right\}$.

It is straightforward to see that such allocations exist: $\delta_1 = \frac{3}{4}$, $\delta_2 = \frac{1}{4}$, $\eta = \frac{2.5}{w}$ is an example. Suppose, for reductio, that a Pareto optimal allocation is m-accessible for some finite m: there exists a sequence of allocations $(x^1_j, x^2_j) = ((\delta_{j1}w, \delta_{j2}w), ((1-\delta_{j1})w, (1-\delta_{j2})w))$, $1 \leqslant j \leqslant m+1$, with $(x^1_1, x^2_1) = (x^1, x^2)$, $x^i_{j+1} \succ^i x^i_j$ or $x^i_{j+1} = x^i_j$ for all i, j, and (x^1_{m+1}, x^2_{m+1}) Pareto optimal – and so $(x^1_{m+1}, x^2_{m+1}) = ((\delta' w, \delta' w), ((1-\delta')w, (1-\delta')w))$ for some $\delta' \in [0, 1]$. Without loss of generality, it can be assumed that $\delta' \leqslant \delta_{(j+1)1} \leqslant \delta_{j1} \leqslant \delta_1$ and $\delta' \geqslant \delta_{(j+1)2} \geqslant \delta_{j2} \geqslant \delta_2$ for all $1 \leqslant j \leqslant m$. Moreover, such a sequence implies that $0.5u^2((1-\delta_1)w) + 0.5u^2((1-\delta_2)w) < u^2((1-\delta')w)$ and $p(s_1)u^1(\delta_1 w) +$

$p(s_2)u^1(\delta_2 w) < u^1(\delta' w)$ for all $p \in \bigcap_{x' \neq x^1} ri(D^1(x^1, x'))$; since $\bigcap_{x' \neq x^1} ri(D^1(x^1, x'))$ contains (only) the probability measure giving the value 0.5 to each state, it follows that $0.5u^1(\delta_1 w) + 0.5u^1(\delta_2 w) < u^1(\delta' w)$. Hence:

(VII.C.3) $$\delta_2 < \delta_1^{0.5}\delta_2^{0.5} < \delta' < 1 - (1-\delta_1)^{0.5}(1-\delta_2)^{0.5} < \delta_1$$

Consider an arbitrary consecutive pair (x_j^1, x_j^2) and (x_{j+1}^1, x_{j+1}^2) in the sequence. By Theorem VII.3 and the fact that the latter is a Pareto-improvement on the former, $\{(\frac{p(s_1)u^{1\prime}(x_j^1(s_1))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))}, \frac{p(s_2)u^{1\prime}(x_j^1(s_2))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))} | \ p \in ri(D^1(x_j^1, x_{j+1}^1))\} \cap \{(\frac{p(s_1)u^{2\prime}(x_j^2(s_1))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))}, \frac{p(s_2)u^{2\prime}(x_j^2(s_2))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))}t) | \ p \in ri(D^2(x_j^2, x_{j+1}^2))\} = \varnothing$. Doing the calculations, and using the fact that $ri(D^2(x_j^2, x_{j+1}^2)) = 0.5$, this is the case if $\frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}} \notin ri(D^1(x_j^1, x_{j+1}^1))$. Hence we must have that:

(VII.C.4) $$\min\{\eta w \max\{|\delta_{j1} - \delta_{(j+1)1}|, |\delta_{(j+1)2} - \delta_{j2}|\}, 0.45\} + 0.5 \leqslant \frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}$$

Note that $0.95 \leqslant \frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}$ if and only if:

$$\begin{aligned} \delta_{j1} - \delta_{j2} &\geqslant 18\delta_{j2}(1 - \delta_{j1}) \\ &\geqslant 18\delta_2(1 - \delta_1) \end{aligned}$$

by the bounds noted above on δ_{j1} and δ_{j2}. It follows from assumption (b) and the fact that $\delta_{j1} - \delta_{j2} \leqslant \delta_1 - \delta_2$ for all j, that, for all j, $0.95 > \frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}$. (VII.C.4) thus reduces to the following inequalities

$$\begin{aligned} \delta_{(j+1)1} &\geqslant \delta_{j1} - \frac{1}{2\eta w}\frac{\delta_{j1} - \delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}} \\ \delta_{(j+1)2} &\leqslant \delta_{j2} + \frac{1}{2\eta w}\frac{\delta_{j1} - \delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}} \end{aligned}$$

And so:

(VII.C.5) $$\delta_{(j+1)1} - \delta_{(j+2)2} \geqslant (\delta_{j1} - \delta_{j2})\left(1 - \frac{1}{\eta w}\frac{1}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}\right)$$

But, using (VII.C.3):

$$\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2} \geqslant \begin{cases} (1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1) & \text{if } \delta_{j1},\delta_{j2} > \frac{1}{2} \\ \delta_1^{0.5}\delta_2^{0.5}+\delta_2-2\delta_1^{0.5}\delta_2^{1.5} & \text{if } \delta_{j1},\delta_{j2} < \frac{1}{2} \\ \frac{1}{2} & \text{if } (\delta_{j1}-\frac{1}{2})(\delta_{j2}-\frac{1}{2}) \leqslant 0 \end{cases}$$

It thus follows from (VII.C.5) that:

(VII.C.6) $$\delta_{(j+1)1}-\delta_{(j+1)2} \geqslant (\delta_{j1}-\delta_{j2})(1-\chi)$$

where

$$\chi = \max\left\{\frac{1}{\eta w}\frac{1}{(1-\delta_1)+(1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1-1)}, \frac{1}{\eta w}\frac{1}{\delta_1^{0.5}\delta_2^{0.5}+\delta_2-2\delta_1^{0.5}\delta_2^{1.5}}, \frac{2}{\eta w}\right\}$$

By assumption (c), $\chi < 1$. Iterating inequality (VII.C.6), we obtain:

$$\delta_{(m+1)1}-\delta_{(m+1)2} \geqslant (\delta_1-\delta_2)(1-\chi)^m > 0$$

contradicting the assumption that (x^1_{m+1}, x^2_{m+1}) is Pareto optimal.

□

Bibliography

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Aumann, R. J. (1962). Utility Theory without the Completeness Axiom. *Econometrica*, 30(3):445–462.

Bewley, T. F. (1986 / 2002). Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2):79–110.

Bewley, T. F. (1989). Market Innovation and Entrepeneurship: A Knightian View. Technical Report 905, Cowles Foundation.

Billot, A., Chateauneuf, A., Gilboa, I., and Tallon, J.-M. (2000). Sharing Beliefs: Between Agreeing and Disagreeing. *Econometrica*, 68(3):685–694.

Danan, E. (2003a). A behavioral model of individual welfare. Technical report, Université Paris 1.

Danan, E. (2003b). Revealed Cognitive Preference Theory. Technical report, EUREQua, Université de Paris 1.

Dubra, J., Maccheroni, F., and Ok, E. A. (2004). Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133.

Faro, J. H. (2013). Variational Bewley Preferences. Technical report, Insper.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Fon, V. and Otani, Y. (1979). Classical welfare theorems with non-transitive and non-complete preferences,. *Journal of Economic Theory*, 20(3):409–418.

Galaabaatar, T. and Karni, E. (2013). Subjective expected utility with incomplete preferences. *Econometrica*, 81(1):255–284.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2):133–173.

Ghirardato, P. and Siniscalchi, M. (2014). Risk-Sharing in the Small and in the Large.

Gilboa, I., Maccheroni, F., Marinacci, M., and Schmeidler, D. (2010). Objective and Subjective Rationality in a Multiple Prior Model. *Econometrica*, 78(2):755–770.

Gilboa, I., Postlewaite, A., and Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy*, pages 285–296.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Gul, F. and Pesendorfer, W. (2006). Random expected utility. *Econometrica*, 74(1):121–146.

Hill, B. (2009). Confidence in preferences. *Social Choice and Welfare*, 39(2):273–302.

Hill, B. (2013). Confidence and decision. *Games and Economic Behavior*, 82:675–692.

Hill, B. (2014). Confidence as a Source of Deferral.

Kopylov, I. (2009). Choice deferral and ambiguity aversion. *Theoretical Economics*, 4(2):199–225.

Lehrer, E. and Teper, R. (2011). Justifiable preferences. *Journal of Economic Theory*, 146(2):762–774.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 74(6):1447–1498.

Masatlioglu, Y. and Ok, E. A. (2005). Rational choice with status quo bias. *Journal of Economic Theory*, 121(1):1–29.

Minardi, S. and Savochkin, A. (2013). Preferences With Grades of Indecisiveness. Technical report, HEC Paris.

Nau, R. (2006). The shape of incomplete preferences. *The Annals of Statistics*, 34(5):2430–2448.

Nau, R. F. (1992). Indeterminate Probabilities on Finite Sets. *The Annals of Statistics*, 20(4):1737–1767.

Nehring, K. (2009). Imprecise probabilistic beliefs as a context for decision-making under ambiguity. *Journal of Economic Theory*, 144(3):1054–1091.

Ok, E. A., Ortoleva, P., and Riella, G. (2012). Incomplete Preferences Under Uncertainty: Indecisiveness in Beliefs versus Tastes. *Econometrica*, 80(4):1791–1808.

Rigotti, L. and Shannon, C. (2005). Uncertainty and Risk in Financial Markets. *Econometrica*, 73(1):203–243.

Rigotti, L., Shannon, C., and Strzalecki, T. (2008). Subjective Beliefs and Ex Ante Trade. *Econometrica*, 76(5):1167–1190.

Schmeidler, D. (1971). A Condition for the Completeness of Partial Preference Relations. *Econometrica*, 39(2):403–404.

Seidenfeld, T., Schervish, M., and Kadane, J. (1995). A representation of partially ordered preferences. *The Annals of Statistics*, 23(6):2168–2217.

VIII Confidence as a Source of Deferral

Abstract

One apparent reason for deferring a decision – abstaining from choosing, leaving the decision open to be taken by someone else, one's later self, or nature – is for lack of sufficient confidence in the relevant beliefs. This paper develops an axiomatic theory of decision in situations where a costly deferral option is available that captures this source of deferral. Drawing on it, a preliminary behavioural comparison with other accounts of deferral, such as those based on information asymmetry, is undertaken, and a simple multi-factor model of deferral – involving both confidence and information considerations – is formulated. The model suggests that incorporating confidence can account for cases of deferral that traditional accounts have trouble explaining. [1]

Keywords: Confidence; multiple priors; deferral; delegation; information acquisition; value of information; incomplete preferences.

JEL classification: D81, D80, D83.

1. This is a slightly modified version of HEC Paris Research Paper No. ECO/SCD-2014-1060.

VIII.1 Introduction

VIII.1.1 Motivation

'If you're not sure, say so.' At first glance, this appears to be a reasonable guide to when a person should *defer*: that is, when he should abstain from choosing any of the options on offer, leaving the decision open, to be taken by someone else, by his future self, or by nature. If you're not sure about any of the options on offer, deferring is one way of saying so. However, this maxim is largely ignored by standard economic accounts of deferral.

Deferral or delegation of decisions to another agent is normally analysed in terms of the expected difference in information (or information-gathering ability) between the 'deferrer' and the 'deferree', without taking into account how sure the former would be in taking the decision himself (for example, Holmstrom (1984)). One family of theories of deferral to one's future self is based on expectations about what information is available to one's future self or what his preferences are (for example, Stigler (1961); Marschak and Miyasawa (1968); Koopmans (1964); Kreps (1979)). Another family considers deciding to be a task, and analyzes deferral – or procrastination of that task – in terms of the comparison between the value of choosing now and the discounted value of choosing at a future moment, in the presence of time inconsistencies or salience effects (for example, Akerlof (1991); O'Donoghue and Rabin (2001)). Under all of these accounts, whether the decision maker is sure or not about which is the best option is assumed not to be a factor in his deferring the decision. The sources of deferral involved are *extrinsic*: they inevitably refer to factors beyond the immediate decision under consideration (the information held by another agent or by one's future self, one's future preferences, the value of accomplishing the task of deciding tomorrow). By contrast, not being sure about what to choose is a source of deferral that only makes reference to the decision maker's own attitudes at the moment of decision: it is *intrinsic*.

Despite their absence from accounts of deferral to others or to one's future self, intrinsic considerations do seem to be involved in some theories of status quo choice – which can be thought of as an instance of deferral to nature. Most prominent are those involving incomplete preferences and according to which the decision maker defers to the status quo in the absence of appropriate determinate preference (for example, Bewley (2002)) . To the extent that his preferences are incomplete, he is not sure of the best option, and this plays a role in his deferral to the status quo.

Beyond the plausible intuition that how sure a decision maker is may play a role in his deferral behaviour, the peculiarity of the current state of the literature – where some of our best models of deferral to nature involve intrinsic sources, whilst the main models of deferral to others or to one's future self assume that such factors have no role to play – pleads in favour of the development of theories of intrinsic deferral general enough to cover the latter sorts of cases.

Moreover, intrinsic sources need not be without significant economic consequences, even for deferral to others or to one's future self. Consider, for example, a decision maker faced with an investment decision that can be deferred (or delegated) to a portfolio manager, and suppose that the information differential between the decision maker and the portfolio manager does not justify the manager's fee; mutual fund data suggests that this situation may be quite realistic (Malkiel, 1995; Gruber, 1996). Whilst the information differential and the cost of delegation cannot explain the (observed) deferral in such cases, incorporation of intrinsic sources of deferral can: the decision maker defers because he is not sure how to invest. Or, to take another example, consider a single, important investment decision, among several given options, where all the information one could expect to learn in a reasonable timeframe is already available, and suppose that the decision maker may defer the decision to his future self (or procrastinate); the choice of whether and how to invest in a 401(k) pension plan may be a decision of this sort. Given the absence of a significant information advantage to deferring, and given that, under the time inconsistency-salience approach, there is a lower tendency to procrastinate on single, important tasks – such as deciding one's investment policy for retirement – it may not be straightforward to account for deferral in such cases on the basis of extrinsic sources alone.[2] Once again, incorporation of intrinsic sources of deferral – most people are not very confident in their judgements about the best investment policies – could contribute to understanding deferral of these sorts of decisions. Similarly, deciding on the basis of tealeaves, palm-reading, horoscopes or oracles, or delaying the decision to consult one of these sources, are forms of deferral – to others or to one's future self – for which information acquisition is too scant (even by the

2. O'Donoghue and Rabin (2001) show that procrastination falls as importance rises when there is a single task to be accomplished. To account for procrastination of important decisions, they develop a model of procrastination in the presence of several tasks. For some further discussion of the relationship to their model, see Section VIII.4.2, footnote 26.

lights of many who rely on these methods) and time-consistency considerations too weak to explain, and where the decision maker's lack of confidence about the appropriate choice may have a role to play.[3] As these examples suggest, in natural situations where decision makers are not sure about what to do – and thus may feel uncomfortable deciding – incorporation of this aspect may help understand cases of deferral that are not clearly accounted for by existing extrinsic approaches.

Any theory of intrinsic deferral with a serious claim of covering all species of deferral – to others, to one's future self and to nature – must be able to tackle the issue of the *cost of deferral*. Deferring to others or to one's future self may incur costs, such as costs of delegation or delay. Standard extrinsic accounts accommodate such costs, roughly speaking, by assigning a value to deferral that is 'weighed off' against its cost. But what is the value of deferral in a situation where the decision maker is not sure what to choose? Existing theories of intrinsic deferral, such as those treating status quo choice, offer no answer to this question.

The aim of this paper is to propose a theory of decision in the presence of a costly deferral option according to which deferral is driven by the decision maker's confidence. This theory will provide the basis for a proper evaluation of this intrinsic source of deferral and its consequences. First of all, since it applies in the costly deferral situations that are the focus of most extrinsic theories of deferral, it will allow a behavioural comparison with these approaches, revealing behavioural differences between them. Secondly, it opens the door to the development of multi-factor models of deferral. Of course, in many real-life situations, several factors – both intrinsic and extrinsic – may be in play. It is standard practice to propose and study models based entirely on one factor (say, information acquisition), ignoring others (for example, time consistency). Likewise, our proposal of a model based solely on intrinsic sources of deferral, ignoring extrinsic sources, is intended as a first step. The ultimate aim is to lay the groundwork for the development of models incorporating both intrinsic and extrinsic sources of deferral. As an illustration, a multi-factor model of deferral, incorporating how sure the decision maker is as well as expected differences in information, will be sketched and some consequences explored.

3. Note that the 'entertainment value' of consulting an oracle does not explain why people tend to consult them specifically when they are faced with a difficult decision. Casual observation suggests that this is indeed the case: indeed, some divination methods are to be used, we have been told, only when there is an important choice to be made.

VIII.1.2 Confidence and the Value of Deferral

As noted previously, the central conceptual difficulty for the development of a theory of intrinsic deferral in the presence of a costly deferral option concerns the question of the value of deferral. To fix ideas, let us restrict ourselves to the framework of decision under uncertainty, where preferences are determined by beliefs and utilities. We focus throughout the paper on the role of the decision maker's confidence in his beliefs, tacitly assuming that the decision maker is fully confident in his utilities. The benchmark in this framework is the model of deferral to nature (or status quo choice) proposed by Bewley (2002). It involves a set of probability measures, which can be thought of as representing the beliefs or probability judgements [4] the decision maker is sure of. Moreover, the unanimity decision rule used – there is a preference between acts if the expected utilities of the acts lie in the appropriate relation for all the probability measures in the set – can be thought of as reflecting the maxim with which this paper began: if the decision maker is not sure of the relevant beliefs, then he does not come to a decision. However, as noted above, the model provides no guidance as to how to value the option of deferral.

The solution proposed in this paper starts from the observation that the representation of beliefs by a set of probability measures assumes that being sure is an all-or-nothing affair: either a belief or probability judgement holds for all measures in the set – the decision maker is sure of it – or not. This assumption is unrealistic: one can be more confident of some beliefs than others. As explored in Hill (2013, 2014), the introduction of a graded notion of confidence in beliefs allows the development of more refined theories of decision. In particular, it permits the level of confidence required of a belief for it to play a role in a decision to depend, say, on the importance of the decision. In the case of interest here, such dependence would yield a theory of deferral encompassing the following reasonable maxim: decide when one has sufficient confidence in the relevant beliefs given the importance of the decision, and defer if not. Or, to put it in the terms of the adage with which this paper began: when you're not sure *enough*, say so.

Beyond its role in *driving* deferral, a graded notion of confidence may also have a role to play in *pricing* deferral. For, once introduced, one can talk not only of how confident

4. By 'probability judgement' we mean a statement concerning probabilities, such as 'the probability of event A is greater than 0.3', held by the agent. We use the terms 'belief' and 'probability judgement' interchangeably in the discussion below.

a decision maker is in certain beliefs, but also of how confident he *would need to be* to decide, given the importance of the decision to be made. It is plausible that, among cases where the decision maker lacks sufficient confidence in the relevant beliefs, the option of deferral is perceived as more valuable in decisions that call for more confidence. As the level of confidence appropriate for the decision goes up, the value of not having to decide – the value of deferring – increases as well. Accordingly, a higher cost of deferral is required to induce the decision maker to decide rather than defer.

Formally, we employ the representation of confidence in beliefs by a *centered confidence ranking* – a nested family of sets of probability measures that includes a singleton set – proposed and defended in Hill (2013). Each set in the family corresponds to a level of confidence, and is understood as representing the beliefs held to that level of confidence. The probability measure in the singleton element of the confidence ranking – called the centre – represents all of the decision maker's beliefs, even those in which he has little confidence. Following the cited paper, the decision maker's attitude to choosing on the basis of limited confidence is represented by a *cautiousness coefficient* – a function assigning to each decision a set in the confidence ranking, interpreted as the appropriate level of confidence (in the decision maker's eyes) for the decision. Finally, we introduce a *cost function* that assigns to each level of confidence a real value, understood as the psychological cost of choosing in a decision requiring that level of confidence when one is not sufficiently confident in the relevant beliefs.

Following the ideas mooted above, we develop a model according to which the decision maker's choice from a menu of Anscombe-Aumann acts A where deferral is possible at a (monetary) cost of x is represented as follows. If there is an optimal element according to the unanimity rule à la Bewley applied with the set of priors picked out by the cautiousness coefficient – that is, if he has sufficient confidence in the relevant beliefs given the decision – then he decides. Otherwise his behaviour depends on the cost of deferral x. If it is outweighed by the cost of deciding – given by the cost function – then he defers. If, on the other hand, the cost of deferral outweighs the cost of deciding, he decides, using all of his beliefs, even those in which he is not very confident (that is, using the probability measure in the singleton element of his confidence ranking).

An axiomatic characterisation of this model is proposed, in which all the aforementioned elements are endogenously derived. The essential axiomatic difference with respect to the standard expected utility model (where there is no deferral) is in appropriate weakenings of the basic choice axioms (*grosso modo*, those playing the role of transitivity, completeness and independence in preference-based models) to allow deferral in some circumstances where standard axioms oblige the decision maker to decide. To that extent, the behavioural properties of the model are quite reasonable. Moreover, the model is behaviourally distinguishable (that is, distinguishable on the basis of the choice patterns it licenses) from a popular theory of costly extrinsic deferral based on the expected information acquisition or asymmetry. Finally, it permits comparative statics analyses of the decision maker's attitudes to deferral, which support the interpretations of the various elements of the model suggested above.

The representation result formally shows how confidence can price deferral, by showing that the confidence-related source of deferral generates a 'cost of deciding', which can be elicited in principle from behaviour. This insight opens the door to the development of multi-factor models of deferral. As an illustration, we formulate and briefly investigate an extension of the model of deferral developed here to incorporate information-related considerations. This model naturally yields a simple additive expression determining deferral, combining the standard value of information and the cost of deciding. It can comfortably accommodate the examples given in the Section VIII.1.1, where important decisions are deferred in the absence of clear extrinsic reasons for deferral; in such cases, the confidence-related cost of deciding will generally be large, and could thus 'tip the balance' in favour of deferral.

The framework and basic notions are introduced in Section VIII.2. The representation is given and axiomatised in Section VIII.3. Section VIII.4 contains a comparison with a standard information-based model of deferral and a brief exploration of a multi-factor model. Related literature not treated elsewhere is discussed in Section VIII.5. Proofs, as well as a comparative statics analysis, are to be found in the Appendices.

VIII.2 Preliminaries

VIII.2.1 Setup

Throughout the paper, we use a version of the Anscombe and Aumann (1963) framework. Let X, the set of outcomes, be a separable metric space. To simplify the treatment of the cost of deferral, we suppose that outcomes are monetary, and hence take $X = \Re$ (with the standard metric). A *consequence* is a probability measure on X with finite support. $\Delta(X)$ is the set of consequences, endowed with the weak topology. Since the space of Borel probabilities measures over X is metrizable (Billingsley, 2009, p72; Aliprantis and Border, 2007, Theorem 15.12), so is $\Delta(X)$. The operation $\dot{-} : \Delta(X) \times X \to \Delta(X)$ is defined as follows: for all $c \in \Delta(X)$ and $x \in X$, $(c \dot{-} x)(y) = c(y + x)$ for all $y \in X$. $c \dot{-} x$ is the result of 'subtracting' a monetary value x from (each outcome yielded by) a consequence c. A function $u : \Delta(X) \to \Re$ is *zeroed* if $u(0) = 0$ and it is *strictly increasing* if, for all $x, y \in X = \Re$, $x \geqslant y$ if and only if $u(x) \geqslant u(y)$.

Let S be a non-empty finite set of states; subsets of S are called *events*. $\Delta(S)$ is the set of probability measures on S, endowed with the Euclidean topology. The objects of choice are *acts*, defined to be functions from states to consequences. $\mathcal{A}$ is the set of acts, endowed with the inherited product metric. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in [0, 1]$, the mixture $\alpha f + (1 - \alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s, x) = \alpha f(s, x) + (1-\alpha)h(s, x)$. We write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$. Similarly, for any $x \in X$, we shall also use x to denote the constant act yielding the degenerate lottery giving x with probability one.

The decision maker is faced with choices from sets of acts, where deferral is possible but costly. [5] To this end, let $\wp(\mathcal{A})$ be the set of non-empty compact subsets of $\mathcal{A}$, endowed with the Hausdorff metric. [6] We call elements of $\wp(\mathcal{A})$ *menus*. Pointwise mixtures of menus are defined as standard: $A_\alpha h = \{f_\alpha h \mid f \in A\}$, for $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$, $\alpha \in [0, 1]$. Let $\dagger_x$ be the option of deferring at (non-negative) cost x, and $\mathcal{D} = \{\dagger_x \mid x \in \Re_{\geqslant 0}\}$. An element $(A, \dagger_x) \in \wp(\mathcal{A}) \times \mathcal{D}$ represents the situation in which the decision maker is called upon to choose from menu A where deferral costs x. A *choice correspondence for costly deferral* is a correspondence $\gamma : \wp(\mathcal{A}) \times \mathcal{D} \rightrightarrows \mathcal{A} \cup \mathcal{D}$ – that is, a function $\gamma : \wp(\mathcal{A}) \times \mathcal{D} \to 2^{\mathcal{A} \cup \mathcal{D}} \backslash \varnothing$ – such that $\gamma(A, \dagger_x) \subseteq A \cup \{\dagger_x\}$. The choice correspondence for costly deferral γ delivers for each situation of the sort described a set containing elements in A – which, as standard,

5. Situations where no deferral option is available can be accommodated in this setup, as limiting cases where the cost of deferral becomes prohibitive.

6. This metric is defined as follows: for $A, B \in \wp(\mathcal{A})$, $\mathbf{h}(A, B) = \max\{\max_{x \in A} \min_{y \in B} d(x, y), \max_{y \in B} \min_{x \in A} d(x, y)\}$, where d is the metric on $\mathcal{A}$.

are interpreted as those which the decision maker is inclined to choose – or $\dagger_x$ – which represents the decision to defer.

Some further notation shall prove useful. Two menus $A, B \in \wp(\mathcal{A})$ are *extensionally equivalent* if there exists a surjective correspondence[7] $\sigma : A \rightrightarrows B$ such that $\gamma(\{f(s), g(s)\}, \dagger_0) = \{f(s), g(s)\}$ for all $f \in A$, $g \in \sigma(f)$ and $s \in S$; if this holds, we write $A \overset{e.e}{\simeq} B$. This is the natural notion of equivalence for menus if, as is standard in formal theories of decision under uncertainty, one treats two acts as essentially the same whenever they yield consequences between which the decision maker is indifferent in every state.

For a utility function $u : \Delta(X) \to \Re$, a set of probability measures $\mathcal{C} \subseteq \Delta(S)$ and a menu $A \in \wp(\mathcal{A})$, let:

(VIII.1)

$$\sup(A, u, \mathcal{C}) = \left\{ f \in A \mid \sum_{s \in S} u(f(s)).p(s) \geqslant \sum_{s \in S} u(g(s)).p(s) \;\; \forall p \in \mathcal{C}, \forall\, g \in \mathcal{A} \right\}$$

$\sup(A, u, \mathcal{C})$ is the set of optimal elements of A – those that are ranked better than all other elements in A – according to the unanimity rule with u and $\mathcal{C}$, which ranks an act better than another if the former has higher expected utility for all probability measures in $\mathcal{C}$. Note that if $\mathcal{C}$ is a singleton, then $\sup(A, u, \mathcal{C})$ is the set of acts in A with maximal expected utility calculated with u and the element in $\mathcal{C}$.

As a final piece of notation, let $\Phi = \{A \in \wp(\mathcal{A}) \mid \exists h \in \mathcal{A},\ \alpha \in (0,1], \dagger_0 \in \gamma(A_\alpha h, \dagger_0)\}$. This is the set of menus such that, for at least one mixture of the menu, the decision maker defers.

VIII.2.2 Confidence ranking and cautiousness coefficient

We adopt with some modification two notions that were introduced in Hill (2013).

Definition VIII.1. A *confidence ranking* Ξ is a nested family of closed, convex subsets of $\Delta(S)$. A confidence ranking Ξ is *continuous* if, for every $\mathcal{C} \in \Xi$, $\mathcal{C} = \overline{\bigcup_{\Xi \ni \mathcal{C}' \subsetneq \mathcal{C}} \mathcal{C}'} = \bigcap_{\Xi \ni \mathcal{C}' \supsetneq \mathcal{C}} \mathcal{C}'$. Ξ is *strict* if, for every $\mathcal{C}_1, \mathcal{C}_2 \in \Xi$ with $\mathcal{C}_1 \subset \mathcal{C}_2$, $(\mathcal{C}_1 \cap (ri(\mathcal{C}_2))^c) \cap ri(\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}) =$

7. A correspondence $\sigma : A \rightrightarrows B$ is surjective if, for all $g \in B$, there exists $f \in A$ with $g \in \sigma(f)$.

$\emptyset$.[8] Ξ is *centered* if it contains a singleton set; in this case, the member of the singleton set is called the *centre* and is denoted p_Ξ.

As mentioned in the Introduction, confidence rankings represent decision makers' beliefs and their confidence in their beliefs. A set in the confidence ranking corresponds to a level of confidence and can be thought of as representing the beliefs held to this level of confidence. For further discussion of the main features of the definition, the reader is referred to Hill (2013). The only new property in the definition is strictness, which implies that the confidence ranking is strictly increasing in a particular sense. It ensures, for example, that as the confidence level increases, the highest value for the probability of an event that is endorsed at that level of confidence also increases, unless this value is already maximal over all confidence levels.

As noted in Hill (2013), a decision maker with a centred confidence ranking is one who, if forced to give his best estimate for the probability of any event, could come up with a single value, although he may not be very confident in it. He is, so to speak, a 'Bayesian with confidence'. Whilst we by no means wish to suggest that all decision makers are of this sort, we focus on decision makers with centred confidence rankings to facilitate the comparison with other approaches to deferral, which generally assume expected utility.[9]

The second notion required is that of a *cautiousness coefficient* for a confidence ranking Ξ, which is defined to be a function $D : \wp(\mathcal{A}) \to \Xi$ satisfying the following three properties.

(Extensionality) For all $A, B, \in \wp(\mathcal{A})$, if $A \overset{e.e.}{\simeq} B$, then $D(A) = D(B)$.

(Continuity) For all $\mathcal{C} \in \Xi$, the sets $\{A \in \wp(\mathcal{A}) | \ D(A) \supseteq \mathcal{C}\}$ and $\{A \in \wp(\mathcal{A}) | \ D(A) \subseteq \mathcal{C}\}$ are closed.

(Φ-Richness) For all $A \in \Phi \subseteq \wp(\mathcal{A})$ and $\mathcal{C} \in \Xi$, there exists $h, h' \in \mathcal{A}$ and $\alpha, \alpha' \in (0, 1]$ such that $D(A_\alpha h) \subseteq \mathcal{C} \subseteq D(A_{\alpha'} h')$.

The cautiousness coefficient can be understood as assigning a level of confidence to a menu: $D(A)$ represents the beliefs held to the level of confidence appropriate for use in the choice from the menu A. As discussed in Hill (2013, 2014), it is a subjective element that

8. For a set X, $\overline{X}$ is the closure of X, and $ri(X)$ is its relative interior.

9. Possible extensions, involving weakening of the centering and strictness properties in particular, are discussed in Remark VIII.2 (Section VIII.3.1).

captures the decision maker's attitude to choosing on the basis of limited confidence. [10] The underlying idea is that the appropriate level of confidence is picked out on the basis of the importance of the decision: for more important decisions, more confidence is required. By contrast with Hill (2013, 2014), who exogenously assume a notion of stakes or importance of a decision, this is left implicit in the notion of cautiousness coefficient used here. [11] Accordingly, several new properties of the cautiousness coefficient are required.

Extensionality states that all that counts in the determination of the level of confidence appropriate for a decision are the values of the consequences of the acts in the menu at the different states. Virtually all formal theories of decision under uncertainty treat extensionally equivalent acts – those that yield consequences between which the decision maker is indifferent in every state – as being essentially the same. Extensionality says that whenever two menus are composed of such acts, the same level of confidence is appropriate for the choice from the menus. Continuity, which is fairly standard, seems reasonable: the level of confidence appropriate for choice from a menu may be altered as the menu changes, but one would not expect it to 'jump' with gradual modifications of the menu.

Φ-Richness is a technical property, which states that the appropriate level of confidence for the choice from a menu can be shifted as far up or down as desired, by considering suitable mixtures of the menu. There is a sense in which (in particular in the presence of an independence axiom; see Section VIII.3.2) the choice from A and the choice from $A_\alpha h$ are the 'same' choice. Nevertheless, these choices need not be of the same importance; accordingly, different levels of confidence may be appropriate for the two choices. To that extent, the latter choice can be thought of as a 'version' of the former choice which calls for the use of beliefs held to the level of confidence appropriate for $A_\alpha h$ rather than A. Φ-Richness simply states that for any choice for which the decision maker defers for some version of the choice (that is, every menu in Φ) and any confidence level, there is a version of the choice, obtained by mixing with an act, for which the appropriate confidence level is above the level in question, and there is a version for which the appropriate confidence level is below that level. The intuition is that mixing with an act can change many of

10. Where the cited papers speak of 'attitude to choosing in the absence of confidence', we prefer 'attitude to choosing on the basis of limited confidence' which strikes us as more in tune with the graded notion of confidence involved.

11. A notion of stakes can however be defined from the notion of cautiousness coefficient used here; see Remark VIII.1 below for further discussion.

the properties of a menu, and in particular the main properties that are relevant for the importance of the choice from it, and for the level of confidence appropriate.

Remark VIII.1. In Hill (2013, 2014), the notion of the stakes involved in a decision was fixed exogenously: in the former paper by stipulating a notion of stakes, and in the latter paper by assuming an exogenously-given stakes relation which orders decisions according to whether the stakes involved in them are higher or lower. The cautiousness coefficient was defined in these papers as a function that respects the stakes relation. The model developed below can be formulated in terms of a stakes relation and a cautiousness coefficient that respects it (in the style of Hill (2014)): it suffices to take D as given and define the stakes relation $\leqq$ on menus by $A \leqq B$ iff $D(A) \subseteq D(B)$. (The stakes relation defined in this way will automatically satisfy versions of the properties given in Hill (2014).) To that extent, the notion of stakes assumed in the cited papers can be thought of as endogenised, or elicited from choice, in the main theorem of this paper (Theorem VIII.1). For further discussion, see Remark VIII.3 (Section VIII.3.3) and Section VIII.5.

VIII.3 A model of costly deferral

We now introduce a representation of choice in the presence of a costly deferral option that captures the role of confidence. The aim in this section is not to offer an exhaustive study of deferral, but rather to propose a relatively simple model of intrinsic deferral that can accommodate cases where deferral is costly, and characterise its behavioural properties.

VIII.3.1 Representation

Consider the following representation of a choice correspondence for costly deferral γ: for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$,

$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \text{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \text{if } \sup(A, u, D(A)) = \varnothing \text{ and } c(D(A)) \leqslant u(x) \\ \{\dagger_x\} & \text{otherwise} \end{cases} \tag{VIII.2}$$

where $u : \Delta(X) \to \Re$ is a von Neumann-Morgenstern utility function, Ξ is a continuous strict, centred confidence ranking, $D : \wp(\mathcal{A}) \to \Xi$ is a cautiousness coefficient and $c : \Xi \to$

$\Re_{\geqslant 0} \cup \{\infty\}$ is a continuous function that is order-preserving and -reflecting with respect to $\subseteq$: that is, $c(\mathcal{C}) \geqslant c(\mathcal{C}')$ if and only if $\mathcal{C} \supseteq \mathcal{C}'$.[12] The function c is called the *cost function*. It can be understood as assigning to any confidence level the psychological cost of deciding – or equivalently the value of deferring – in a choice that calls for that level of confidence but in which one does not hold the relevant beliefs with sufficient confidence to yield a decision. As one would expect, this function is increasing in the confidence level: for decisions requiring more confidence, it is more (psychologically) costly to decide when one would have wanted to defer for lack of sufficient confidence in the appropriate beliefs.

To understand representation (VIII.2), first note that, when deferral is costless, it becomes:

$$\gamma(A, \dagger_0) = \begin{cases} \sup(A, u, D(A)) & \text{if } \sup(A, u, D(A)) \neq \varnothing \\ \{\dagger_0\} & \text{otherwise} \end{cases}$$

Recall that $D(A)$ represents the set of beliefs that the decision maker holds to the level of confidence appropriate for use in the choice from the menu A. $\sup(A, u, D(A))$ contains the acts in A that are better than all other alternatives under the unanimity rule with the set of probability measures $D(A)$. It can be interpreted as containing those acts that the decision maker can conclude to be better than all other acts on the menu, on the basis of the beliefs in which he has sufficient confidence given the choice to be made. Hence, under this representation, the decision maker decides when there is an act that is optimal according to these beliefs. If there is no such act, he defers. When deferral is free, a decision maker represented by (VIII.2) chooses when he has sufficient confidence in the appropriate beliefs given the decision to be made, and defers when he does not.

In the light of this, representation (VIII.2), applied in the general case of choice from a menu A where it costs x to defer, can be understood as follows. If the decision maker has sufficient confidence, his behaviour is the same as in the case where deferral is free: he chooses an optimal act according to the unanimity rule with appropriate beliefs. In the absence of sufficient confidence – when there is no optimal act according to the rule – his behaviour may differ. In these cases, he compares the psychological cost of deciding given the confidence level appropriate for the decision (given, in utility units, by $c(D(A))$) with

12. We take the standard topology on the non-negative extended reals $\Re_{\geqslant 0} \cup \{\infty\}$, namely that produced by the Alexandroff one-point compactification of $\Re_{\geqslant 0}$ (Aliprantis and Border, 2007, Section 2.16).

the cost of deferral in the choice situation in which he is in (which corresponds to the utility of x, $u(x)$). If the monetary cost of deferral imposed on him is lower than the psychological cost of deciding, then he defers, as recommended by the unanimity rule applied with the appropriate confidence level. If the monetary cost of deferral imposed on him outweighs the psychological cost of deciding, then he decides. In these cases, he chooses an act which has highest expected utility calculated with the probability measure in the centre of the confidence ranking. To the extent that this probability measure captures all of the decision maker's beliefs, this amounts to deciding on the basis of all of his beliefs, irrespective of the confidence with which they are held. The intuition here is that, whilst the decision maker should not decide on the basis of a belief in which he has insufficient confidence when he can defer, if deferral is too costly he may as well mobilise all of his beliefs – even those held with little confidence. Given that he is relying on more beliefs, the decision maker may be able to come to a decision; in fact, for decision makers with centered confidence rankings – which, recall, are basically Bayesians with confidence – this will always be the case.

Whilst not the only way of choosing in such situations, the procedure formalised in representation (VIII.2) is certainly not unreasonable. The decision maker will choose whenever he has sufficient confidence in the relevant beliefs. Otherwise, he will allow himself to defer only if the (psychological) cost of deciding outweighs the (monetary) cost of deferral. Finally, when he does not allow himself to defer, he decides on the basis of all of his beliefs, even those in which he has little confidence.

Remark VIII.2. Although we consider a simple rule for choice in the face of costly deferral, it is possible to formulate models, and extend the representation results, to incorporate modifications or refinements.

First of all, one could abandon the assumption that the confidence ranking is centered. In the case where the decision maker has insufficient confidence in the relevant beliefs but decides due to the cost of deferral, one could adopt the same strategy of employing the smallest set of probability measures in the confidence ranking. Since this may be a non-singleton set, some other rule, such as the maxmin EU rule (Gilboa and Schmeidler, 1989), would be required in cases in which the unanimity rule applied on this set does not yield an optimal element.

Secondly, one could abandon the use of the smallest set in the confidence ranking to choose in cases where the decision maker has insufficient confidence but the cost of deferral is high. One possibility is to use the set of probability measures corresponding to the level of confidence appropriate for the choice from the menu, and a decision rule that always yields an optimal act, such as the maxmin EU rule. Another possibility is to use the largest set in the confidence ranking that yields a decision (that is, for which the optimal set of acts under the unanimity rule is non-empty). Whenever the confidence ranking is strict, this latter possibility is equivalent to representation (VIII.2).

Working in a preference framework, Hill (2014) distinguishes and characterises axiomatically ways of deciding based on all of one's beliefs and ways that rely on the beliefs that one holds to the level of confidence appropriate for the decision. Similar techniques could be employed to generalize Theorem VIII.1 to non-centred or non-strict confidence rankings, and to alternative strategies for choosing in the absence of sufficient confidence.

Thirdly, one could imagine using a criterion for deferral different from the one involved in representation (VIII.2), namely $c(D(A)) \leqslant u(x)$. Note that representations with criteria such as $c(D(A)) \leqslant x$ or $c(D(A)) \leqslant -u(-x)$ (where, in the former, c gives the cost in monetary units) are related to representation (VIII.2) by appropriate transformations of the cost function, and so are behaviourally indistinguishable from it. Hence the results below immediately apply to them. Other possible suggestions include taking the difference in cost between the confidence level required for the decision and the highest confidence level at which the decision maker holds sufficient beliefs to decide, or the following criterion based on the difference in the value of deciding after paying the monetary cost of deferral and deciding in the absence of this cost: $c(D(A)) \leqslant \max_{h \in A} \sum_{s \in S} u(h(s)).p_{\Xi}(s) - \max_{h \in A} \sum_{s \in S} u(h(s) \dot{-} x).p_{\Xi}(s)$. Representation theorems for versions of representation (VIII.2) involving these criteria can be developed, using similar techniques to those adopted in the proof of Theorem VIII.1. Finally, one could go beyond monetary outcomes, developing similar representations and results for outcome spaces and costs that are not purely monetary (including, for example, a temporal element).

VIII.3.2 Axioms

To state the axioms, the following definition shall prove useful.

Definition VIII.2. For each $x \in \Re_{\geqslant 0}$, the function $\bar{\gamma}^x : \wp(\mathcal{A}) \to 2^{\mathcal{A}}$ is defined by: $f \in \bar{\gamma}^x(A)$ if and only if, for all $h \in \mathcal{A}$ and $\alpha \in (0,1]$, if $\dagger_x \notin \gamma(A_\alpha h, \dagger_x)$, then $f_\alpha h \in \gamma(A_\alpha h, \dagger_0)$.

To explain the underlying intuition, let us say that the cost x *is not motivating* for a menu A when, for every version of the menu, [13] if the decision maker chooses from it when deferral costs x, then he also chooses when deferral is free. A cost that is not motivating for a menu does not drive any decision taken from it: whenever the decision maker decides at this cost of deferral, he is 'sure enough' to decide even when deferral is costless. By contrast, a cost that is motivating for a menu may drive the decision from it: the decision maker may choose at this cost though he does not choose at some lower cost. So $f \in \bar{\gamma}^x(A)$ says that the cost x is not motivating for A and (the appropriate mixture of) f is among the potential choices from (the corresponding version of) A whenever a decision is made. Note that $\bar{\gamma}^x$ may take as value the empty set; it does so on menus for which x is motivating. Hence $\bar{\gamma}^x$ is not a choice correspondence (which is standardly assumed to take non-empty values).

We now consider several axioms on γ, which are organised into three groups.

VIII.3.2.1 Main Behavioral Axioms

Axiom VIII-A1 (Contraction)**.** For all $A, B \in \wp(\mathcal{A})$ with $A \subseteq B$, all $x \in \Re_{\geqslant 0}$, and all $f \in \mathcal{A}$, if $f \in \bar{\gamma}^x(B)$, then $f \in \bar{\gamma}^x(A)$.

Axiom VIII-A2 (Strong Expansion)**.** For all $A, B \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, $f \in A$ and $g \in A \cap B$, if $f \in \bar{\gamma}^x(A)$ and $g \in \bar{\gamma}^x(B)$, then $f \in \bar{\gamma}^x(A \cup B)$.

Axiom VIII-A3 (Independence)**.** For all $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$ and $\alpha \in (0,1]$ and all $x, x' \in \Re_{\geqslant 0}$ such that $\dagger_x \notin \gamma(A, \dagger_x)$ and $\dagger_{x'} \notin \gamma(A_\alpha h, \dagger_{x'})$, $\gamma(A, \dagger_x)_\alpha h = \gamma(A_\alpha h, \dagger_{x'})$.

Axiom VIII-A4 (Consistency)**.** For all $A \in \wp(\mathcal{A})$, and all $x, x' \in \Re_{\geqslant 0}$ with $x \leqslant x'$, if $\dagger_x \notin \gamma(A, \dagger_x)$ then $\gamma(A, \dagger_{x'}) = \gamma(A, \dagger_x)$.

Axiom VIII-A5 (Centering)**.** For all $A \in \wp(\mathcal{A})$, there exists $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$.

13. Recall from the discussion in Section VIII.2.2 that different versions of a menu are obtained by mixing the menu with an act.

Contraction (VIII-A1) is just Sen's axiom α applied to $\bar{\gamma}^x$; as such, it is well known in the literature as a standard ingredient in the revealed preference theory of complete preferences. Similarly, Strong Expansion (VIII-A2) is the property π introduced by Hill (2009), formulated for $\bar{\gamma}^x$. As discussed in Hill (2009), π can be thought of as the equivalent of Sen's axiom β for incomplete preferences. Firstly, the intuition supporting it is similar to that supporting the standard axiom β. Secondly, it implies β, and is in fact equivalent to it on ordinary choice correspondences. Finally, α and π basically characterise the rationalisation of a generalisation of standard choice correspondences by reflexive, transitive but not necessarily complete binary relations. Lest it help the reader get a grasp on these axioms, one can think of α and π as boiling down, in the case of binary menus (and hence standard preferences), to the assumption of reflexivity and transitivity without completeness.

To appreciate the formulation of these conditions for $\bar{\gamma}^x$, it is useful to consider the application of Strong Expansion (VIII-A2) to menus $\{f, g\}$, $\{g, h\}$ and $\{f, g, h\}$.[14] It implies in particular that, if f is chosen over g, g is chosen over h, and the cost x is not motivating for these decisions – that is, whenever the decision maker chooses at cost x, he also chooses when deferral is free – then this cost cannot be motivating for the choice from $\{f, g, h\}$ – if he chooses at cost x, then he must also choose when deferral is free. The underlying intuition is reasonable: if the decisions taken between f and g and between g and h at cost of deferral x are driven not by the cost of deferral, but say by the decision maker being 'sure enough', then the same holds for the choice from $\{f, g, h\}$.

By contrast, consider the following formulation of π for choice behaviour when deferral is free, which is arguably a more standard version of the property in the current framework:

Axiom VIII-A2' (Strong Expansion$_{Free}$)**.** For all $A, B \in \wp(\mathcal{A})$, $f \in A$ and $g \in A \cap B$, if $f \in \gamma(A, \dagger_0)$ and $g \in \gamma(B, \dagger_0)$, then $f \in \gamma(A \cup B, \dagger_0)$.

Applied to the case discussed above, Strong Expansion$_{Free}$ (VIII-A2') implies that, if the decision maker chooses f over g and g over h when deferral is free, then he cannot defer from $\{f, g, h\}$ when deferral is free. This may be unreasonably strong in some circumstances. For example, suppose that g represents taking out \$100 of credit to buy a financial product, where the current portfolio is given by h, and f itself represents buying another \$100 product on credit with respect to current portfolio g. Whilst one may be able

14. This case has the advantage of highlighting the relation to the transitivity axiom on preferences. Similar considerations apply to Contraction (VIII-A1).

to decide in each of the binary choices between f and g and g and h, when the credit under consideration is larger (up to \$200 over one's current portfolio) the decision may be more difficult, and it might not be unreasonable to defer.[15] Strong Expansion (VIII-A2) can comfortably accommodate such behavioural patterns: it allows deferral in such cases, albeit under certain conditions. In particular, it only prohibits deferral if the decision maker chooses from $\{f, g, h\}$ at a cost of deferral that is not motivating for the other choices (because VIII-A2 implies that this cost is not motivating for the choice from $\{f, g, h\}$ either).

Note moreover that, whenever a choice is made from $\{f, g, h\}$, then VIII-A2 agrees with VIII-A2' that f will be among the chosen options. The only 'violations' of the more standard Strong Expansion$_{Free}$ that are sanctioned by Strong Expansion are not forms of inconsistent choices (such as those corresponding to preference cycles), but cases of deferral where the standard axiom would have demanded decision. Summing up in terms of preferences for readers more familiar with these: on binary menus, Contraction and Strong Expansion basically amount to the assumption of reflexivity and a weakening of transitivity to allow indeterminacy in some cases where the standard axiom demands determinate preference.

Independence (VIII-A3) demands that the standard independence axiom, formulated in a choice-theoretic setting, holds whenever the decision maker decides rather than deferring. Evidently it fully retains the intuitions behind the standard axiom, and is equivalent to it in the case where the decision maker never defers. The restriction to pairs of menus from which the decision maker decides is a central behavioural difference between the proposed model and more standard ones. Indeed, in the presence of the other axioms, dropping the condition that $\dagger_{x'} \notin \gamma(A_\alpha h, \dagger_{x'})$ from VIII-A3 yields the standard expected utility representation (formulated in the current framework), under which the decision maker never defers. In the current context, this condition is entirely natural: without it, the axiom would demand that if one chooses from A when deferral costs \$100 then one cannot defer from $A_\alpha h$ when deferral is free.

Consistency (VIII-A4), requires an immediately intuitive relationship between decisions with different costs of deferral. It says that, if one is willing to choose from a menu at a particular cost of deferral, then as deferral becomes more expensive, one will continue

15. Consider also the version of this example that involves the successive application of Strong Expansion$_{Free}$ to 100 binary choices in which one decides between taking \$100 credit over one's current portfolio or not.

to choose, and make the same choice. Centering (VIII-A5) characterises the centeredness property of the confidence ranking (Section VIII.2.2). Recall from the discussion in Section VIII.2.2 that different versions of a choice can be obtained by mixing the menu with an act; these different versions may be of varying importance, with different levels of confidence appropriate. In the light of this, Centering states that, for any menu, there is a version of the choice from the menu in which the decision maker will decide rather than deferring. The basic intuition is that, if the appropriate level of confidence is sufficiently low, the decision maker will hazard a choice, albeit one in which he may not be very confident. As mentioned in Section VIII.2.2, the centeredness property of confidence rankings is adopted to ease comparison with standard approaches to deferral; the same goes for the Centering axiom, which will be dropped in any extension of the model that does not involve centered confidence rankings.

VIII.3.2.2 Remaining Behavioral Axioms

Axiom VIII-A6 (Defer or Choose)**.** For all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$, if $\dagger_x \in \gamma(A, \dagger_x)$, then $\gamma(A, \dagger_x) = \{\dagger_x\}$.

Axiom VIII-A7 (No Deferral)**.** For all $A \in \wp(\mathcal{A})$ with $A \subseteq \Delta(X)$ and all $x \in \Re_{\geqslant 0}$, $\dagger_x \notin \gamma(A, \dagger_x)$.

Axiom VIII-A8 (Monotonicity)**.** For all $f, g \in \mathcal{A}$, $z, z' \in X$, $A, B \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$:

i. if $z > z'$ then $\gamma(\{z, z'\}, \dagger_0) = \{z\}$;

ii. if $g(s) \in \gamma(\{f(s), g(s)\}, \dagger_0)$ for all $s \in S$, then $g \in \gamma(\{f, g\}, \dagger_0)$;

iii. if $A \overset{e.e}{\simeq} B$, then $\dagger_x \in \gamma(A, \dagger_x)$ if and only if $\dagger_x \in \gamma(B, \dagger_x)$.

Defer or Choose (VIII-A6), which is specific to the sort of choice situation under consideration here, basically says that if the decision maker defers – $\dagger_x$ is in the choice set – then he does not choose – no acts are retained as admissible choices. No Deferral (VIII-A7) simply states that in choices among constant acts only, the decision maker never defers. The axiom translates the fact that the agent is assumed to be fully confident in his utilities; as stated in the Introduction, only confidence in beliefs is at issue here.

The first clause of Monotonicity (VIII-A8) is a standard assumption in the context of monetary outcomes, saying that the decision maker chooses (strictly) more money over less. It is retained to simplify the representation (insofar as it implies that different costs of deferral are perceived as different by the decision maker), and can be dropped without significant changes to the basic form of the result below. The second clause is a choice-theoretic version of the standard monotonicity axiom for decision under uncertainty as formulated in the preference framework. The final clause demands that whether the decision maker defers at a given cost of deferral respects extensional equivalence of menus: for any two menus such that each act in one is extensionally equivalent to an act in the other (they yield consequences between which the decision maker is indifferent in all states), the decision maker defers from one menu at a given cost of deferral precisely when he defers from the other menu at this cost. Virtually all formal theories of decision under uncertainty treat extensionally equivalent acts in the same way; this axiom extends this to the case of deferral. (The extension to the case of deferral is not a consequence of clause VIII-A8 part ii., as is standardly the case, because of the weakness of the choice-theoretic axioms.)

VIII.3.2.3 Technical Axioms

Before stating the final, technical, axioms, let us introduce the following terminology. Define $\iota : \wp(\mathcal{A}) \rightarrow \Re_{\geqslant 0} \cup \{\infty\}$ by $\iota(A) = \inf\{x \in \Re_{\geqslant 0} \mid \dagger_x \notin \gamma(A, \dagger_x)\}$,[16] and $mmc : \wp(\mathcal{A}) \rightarrow \Re_{\geqslant 0} \cup \{\infty\}$ by $mmc(A) = \inf\{x \in \Re_{\geqslant 0} | \ \bar{\gamma}^x(A) = \varnothing\}$. Moreover, let $\Upsilon = \{A \in \wp(\mathcal{A}) \mid \dagger_0 \in \gamma(A, \dagger_0)\}$.

Υ is the set of menus from which the decision maker defers when deferral is free. For such menus, the cost of deferral counts: whether the decision maker chooses or defers depends on it. ι gives the point where the cost of deferral 'bites': whenever the cost is below the value given by ι the decision maker defers, whereas above that value he decides. Naturally, if a cost is above the 'biting point' for such a menu, then it is motivating for the menu, in the sense introduced previously. The interpretation of ι as a biting point only holds on Υ (ι takes the value zero elsewhere), but mmc can be thought of as an extension beyond this set. It gives the lowest cost that can be motivating,[17] and hence the lowest biting point over all versions of the menu where the cost counts. For any menu A whose biting point is above $mmc(A)$, this point is given by $\iota(A)$, and A is in Υ. So for a menu A not in Υ, although the biting point is not pinned down by $\iota(A)$, one can nevertheless

16. Throughout, we take an infimum to be infinite when the set over which it is taken is empty.

17. Hence the terminology: mmc stands (albeit loosely) for minimal motivating cost. Recall from the remarks following Definition VIII.2 that $\bar{\gamma}^x(A) = \varnothing$ when x is motivating for A.

conclude that it is not greater than $mmc(A)$. Note that Φ (defined in Section VIII.2.1) is the set of menus for which some cost is motivating.

Now consider the following axioms.

Axiom VIII-A9 (Continuity)**.** For all $A, A_n \in \wp(\mathcal{A})$, $f, f_n, g, h \in \mathcal{A}$, $\alpha \in (0,1)$, $x, x_n \in \Re_{\geqslant 0}$:

i. if $A_n \to A$, $f_n \to f$ and $x_n \to x$, and $f_n \in \gamma(A_n, \dagger_{x_n})$ for all $n \in \mathbb{N}$, then $f \in \gamma(A, \dagger_x)$;

ii. if $\sup \iota(\Upsilon) > mmc(\{f, g\})$, $f \in \gamma(\{f, g\}, \dagger_0)$, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, then there exists $y \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^y(\{f, g\}) = \varnothing$ but $\bar{\gamma}^y(\{f, g_\alpha h\}) \neq \varnothing$;

iii. if $x \in (0, \sup \iota(\Upsilon))$, $(\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)$ and $\iota^{-1}([x, \infty]) \cup (mmc^{-1}([x, \infty]) \cap \Upsilon^c)$ are closed in Φ;

iv. $\{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A) = \varnothing\} = \{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A_\alpha h) = \varnothing\}$ and the set is open in $\Re_{\geqslant 0}$.

Axiom VIII-A10 (Richness)**.** For all $A \in \Phi$:

i. for all $x \in \iota(\Upsilon)$, there exist $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $\dagger_y \in \gamma(A_\alpha h, \dagger_y)$ for all $y < x$;

ii. if there exists $x > \inf \iota(\Upsilon)$ with $\bar{\gamma}^x(A) \neq \varnothing$, then $\{A_\alpha h \in \Phi \mid \alpha \in (0, 1], h \in \mathcal{A}\} \nsubseteq \overline{\Upsilon}$.

As indicated, these are basically technical axioms. The first clause of Continuity (VIII-A9) is a version of the standard upper hemi-continuity property for correspondences, with the addition of continuity under changes in the cost of deferral. It captures essentially the same intuition that small changes in the choice situation – in the menu and the cost of deferral – do not induce large changes in choice behaviour. The second clause (VIII-A9 part ii.) states that, for pairs of acts f and g where f is chosen over g when deferral is free, whenever g is 'worsened' – by mixing it with an act dominated statewise by g and f – some costs cease to be motivating. As g is 'worsened' the choice becomes 'easier', so some costs that could drive the initial decision – in the sense that the decision maker is not sure enough to decide when deferral is free – no longer drive the 'easier' decision – he is sure enough in that decision to decide when deferral is free. The third clause (VIII-A9 part iii.) requires a certain continuity in the point at which the cost of deferral bites: as one gradually changes the menu, its biting point (as given by the value of ι if the cost of deferral counts, and the lowest motivating cost mmc if not) does not suddenly 'jump'. The final clause (VIII-A9 part iv.) demands that the set of costs that are motivating for a menu is the same for the menu and any version of it, and that this set is open. The assumption of openness, like some of the other parts of the technical axioms, is largely conventional: appropriate modifications of it would lead, for example, to a representation with the same general form as (VIII.2) but with a strict rather than weak inequality in the condition on the second line.

Richness (VIII-A10) is also a technical axiom. For any menu for which some cost is motivating, the first clause (VIII-A10 part i.) demands that the point at which cost of deferral bites can be moved as high up the range of possible values as desired by taking the appropriate version of the menu. The second clause (VIII-A10 part ii.) states that either every cost is motivating for the menu – and hence the biting point can be moved as low down as desired by considering different versions of the menu – or a condition is satisfied. The condition basically states that for some versions of the menu, the biting point is not entirely fixed by deferral behaviour: the cost does not count for the choice from them (and they are not arbitrary close to menus for which the cost counts), so ι does not give the biting point. In other words, VIII-A10 part ii. says that, if it cannot be concluded that the biting point can be moved as low down the range of possible values as desired (there is a cost that is not motivating for the menu), then it cannot be concluded that it *cannot* be moved as low down the range of values as desired (for some versions of the menu, the cost does not count, so the biting point cannot be fully pinned down).

VIII.3.3 Result

We have the following representation result.

Theorem VIII.1. *Let γ be a choice correspondence for costly deferral on $\mathcal{A}$. The following are equivalent:*

(i) *γ satisfies VIII-A1–VIII-A10,*

(ii) *there exists a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \rightarrow \Re$, a continuous strict centred confidence ranking Ξ with centre p_Ξ, a cautiousness coefficient $D : \wp(\mathcal{A}) \rightarrow \Xi$ for Ξ, and a cost function $c : \Xi \rightarrow \Re_{\geqslant 0} \cup \{\infty\}$, such that,*

for all $A \in \wp(\mathcal{A})$ *and* $x \in \Re_{\geqslant 0}$*:*

(VIII.2)
$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \textit{if } \sup(A, u, D(A)) = \varnothing \textit{ and } c(D(A)) \leqslant u(x) \\ \{\dagger_x\} & \textit{otherwise} \end{cases}$$

Furthermore, Ξ *is unique, the restriction of* D *to* $\overline{\Upsilon}$ *is unique, and* u *and* c *are unique up to the same positive affine transformation.*

The confidence ranking, the utility function and the cost function have the standard uniqueness properties. Moreover, the cautiousness coefficient has the sort of uniqueness properties one would expect. Wherever the confidence level appropriate for the choice from a menu counts – wherever there are some costs of deferral for which the decision maker defers and others where he decides – the level of confidence is fixed uniquely by the cautiousness coefficient. Wherever the decision maker chooses no matter the cost of deferral, the precise value of the appropriate level of confidence is not important; all that matters is that it is low enough for the beliefs held to that level of confidence to yield a choice. On these menus, where the precise setting of the appropriate level of confidence does not matter for choice, the cautiousness coefficient is not necessarily unique. In other words, the cautiousness coefficient is unique where the appropriate level of confidence matters, but not where it doesn't.

Remark VIII.3. As noted in Remark VIII.1, the representation can be reformulated in terms of a stakes relation on the set of menus and a cautiousness coefficient that respects this relation, in the style of Hill (2014). This reformulation gives a version of Theorem VIII.1 that yields a stakes relation, as well as the other elements cited. It is straightforward to check that this stakes relation has the same uniqueness properties as the cautiousness coefficient: it is unique on $\overline{\Upsilon}$. To this extent, Theorem VIII.1 provides an independent contribution to the literature on the confidence-based approach set out in Hill (2013, 2014), in the form of a representation result where the stakes relation is endogenised, that is elicited from behaviour rather than assumed. See Section VIII.5 for further discussion.

We note finally that comparative statics analyses can be undertaken on the proposed representation, separating in particular the roles of various elements of the model. Some details are given in Appendix VIII.A.

VIII.4 Confidence and difference in information as sources of deferral

The representation and characterization of choice in the presence of a costly deferral option lays the ground for a deeper understanding of the relationship between the intrinsic source of deferral explored in this paper and extrinsic sources studied elsewhere in the literature (Section VIII.1.1), as well as for the development of multi-factor theories of deferral. As a preliminary exploration, we focus on a particular, popular model of extrinsic deferral, based on the difference in information between the decision maker and the person who will finally take the decision. We perform a brief comparison with the model proposed in the previous section, and consider how the two aspects could be combined.

VIII.4.1 Confidence- versus information-driven deferral

Assume the setup used above, and suppose moreover that if the decision maker – call him DM1 – defers, then the person to whom he defers (who may be his future self) – call him DM2 – will be faced with the same choice, which he cannot defer. A standard simple approach to this sort of decision would be to assume that both decision makers are Bayesian with the same (strictly increasing, zeroed, continuous, affine) utility function $u : \Delta(X) \to \Re$.[18] The information differential between the decision makers can be characterised by a finite set I of signals, with each $i \in I$ associated with an ex post probability measure p_i over S, and a probability measure q over I. DM2 receives one of the signals before deciding, and so has probability measure p_i for the appropriate i. DM1 only has a probability measure over the signal received by DM2, q; his probability measure over states is p where $p(s) = \mathbb{E}_q p_i(s)$ for all $s \in S$.[19] The value of deferral, measured in utility units, is the value of the expected information differential given the cost: that is, the difference between the expected value of choosing ex post after having paid the cost of deferral and the value of choosing ex ante. More precisely, for the choice from a menu A where deferral costs x, it

18. The strategic literature on delegation only gets properly started, of course, when the utility functions are different, but we focus on this simple case to highlight the contrast with the model proposed here.

19. $\mathbb{E}_q$ is the expectation operator given distribution q, defined by $\mathbb{E}_q f(i) = \sum_{i \in I} f(i) dq(i)$.

is given by: [20]

$$VI_x(A) = \mathbb{E}_q \max_{h \in A} \mathbb{E}_{p_i} u(h(s) \dot{-} x) - \max_{h \in A} \mathbb{E}_p u(h(s)) \tag{VIII.3}$$

A formulation in monetary units is given by the demand value of deferral (Marschak and Miyasawa, 1968) – the highest monetary amount that DM1 is willing to pay to defer – which, for a menu A, is the (unique) solution z to

$$VI_z(A) = 0 \tag{VIII.4}$$

The decision maker (DM1) defers whenever $VI_x(A)$ is positive, or equivalently the cost of deferral is less than the demand value, and decides if not. Hence this approach yields the following representation of the choice correspondence for costly deferral γ: for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$

$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, \{p\}) & \text{if } VI_x(A) \leqslant 0 \\ \{\dagger_x\} & \text{otherwise} \end{cases} \tag{VIII.5}$$

As concerns the criteria determining whether the decision maker defers or not, there are two evident points of contrast with representation (VIII.2). The first concerns the value or cost considerations driving deferral. Under the standard approach, it is the value of the expected information that determines whether deferral is acceptable, whereas under the model proposed in Section VIII.3, the relevant factor is the (psychological) cost of deciding when one would have wanted to defer. Note that, consistently with the distinction between extrinsic and intrinsic reasons for deferring introduced in Section VIII.1.1, the value of information depends on factors beyond DM1's immediate decision – and in particular the information available to the person who will take the decision if it is deferred – whilst the cost of deciding in representation (VIII.2) only depends on the decision itself, and in particular the level of confidence appropriate for it. Given the relative nature of the former term, this difference between the two representations does not immediately translate into a perspicuous axiomatic difference in the setup adopted here.

20. Recall from Section VIII.2.1 that $h(s) \dot{-} x$ is the result of subtracting the monetary value x from each outcome yielded by the lottery $h(s)$.

The second point of contrast concerns the existence of a second 'non-cost' criterion determining whether the decision maker defers or not. Under the value of information approach, deferral is entirely dictated by the value considerations just discussed: the decision maker defers whenever the value of the expected information given the cost of deferral is positive. Under representation (VIII.2), the decision maker's capacity to decide on the basis of the beliefs in which he has appropriate confidence also plays a role in determining whether he defers or not. In particular, he does not necessarily defer in cases where the psychological cost of deciding outweighs the cost of deferral: whenever he has sufficient confidence in the relevant beliefs, he decides. This difference between the value of information and confidence approaches does have simple behavioral consequences.

Under the value of information approach, if the decision maker defers the decision from a menu A at some non-zero cost of deferral, then for each mixture $A_\alpha h$ of the menu, there is a non-zero cost of deferral at which the decision maker defers the choice. For instance, if the decision maker defers, at some cost, the choice of whether to take a bet where one could win or lose \$1 M, then there is a cost he would pay to defer the choice of whether to bet on the same event but with stakes of \$10.[21] If there is a difference of information between DM1 and DM2 in one case, there will be an information difference in the other. This behavioural pattern may be violated under the model developed in this paper, because, although one might not have enough confidence in the relevant beliefs to decide when \$1 M are at stake, the level of confidence required may be lower when only \$10 are at stake. If one holds the relevant beliefs to that level of confidence, one will decide in the latter decision, no matter the cost of deferral.[22]

We thus conclude that the confidence approach developed in this paper is behaviorally distinguishable from one of the major standard approaches to deferral, or delegation, namely that based on difference or asymmetry of information.[23]

21. The example in the text assumes a linear utility function. As standard, this assumption can be dropped by replacing \$10 by a lottery yielding a 10^{-5} chance of getting \$1 M.

22. More formally, over and above the independence axiom used here (VIII-A3), the information-driven model of deferral (VIII.5) satisfies the following condition, whilst the confidence-based model (VIII.2) does not:

(Deferral-Independence). For all $A \in \wp(\mathcal{A})$, $h \in \mathcal{A}$, $\alpha \in (0, 1]$ and $x \in \Re_{>0}$, if $\dagger_x \in \gamma(A, \dagger_x)$, then there exists $x' \in \Re_{>0}$ such that $\dagger_{x'} \in \gamma(A_\alpha h, \dagger_{x'})$.

23. A similar analysis applies *mutatis mutandis* to many of the main models in which deferral (leaving a choice open) is based on consideration of the decision maker's future self's possible utility functions (for

VIII.4.2 Combining confidence and information: a simple multi-factor model of deferral

As made clear at the outset, the ultimate aim of this paper is not to claim that there is one 'correct' account of deferral, but rather to draw attention to a different, intrinsic, source of deferral, which eventually may be integrated into existing approaches. To illustrate its potential interest, we now consider, on a very simple setup, how the two reasons for deferring discussed above – expected difference in information and insufficient confidence – may be combined, and some consequences of the presence of the two.

To this end, consider a decision maker DM3 who perceives the same information structure as DM1 (Section VIII.4.1), but who is also a confidence-based decision maker of the sort proposed in Section VIII.3. More specifically, suppose that, beyond the utility function, the set of signals I, posteriors p_i and the distribution q over I, the decision maker has a centered confidence ranking that is centered on p (DM1's prior probability measure).[24] Whenever DM3 has sufficient confidence given the decision, he chooses, as in representation (VIII.2). In these cases, he makes the same choices as DM1. When he lacks sufficient confidence, his behaviour will depend on the value of deferring, which now incorporates confidence- and information-related factors. The value, in utility units, of deferring from menu A when deferral costs x and the decision maker has insufficient confidence in the relevant beliefs becomes:

$$\text{(VIII.6)}\quad \begin{aligned} VD_x(A) &= \mathbb{E}_q \max_{h\in A} \mathbb{E}_{p_i} u(h(s) \dot{-} x) - \max_{h\in A} \mathbb{E}_p \left(u(h(s)) - c(D(A))\right) \\ &= VI_x(A) + c(D(A)) \end{aligned}$$

example, Kreps (1979); Dekel et al. (2001); Gul and Pesendorfer (2001)). Note in particular that the addition of a self-control cost to a representation involving difference in information (including information about future preferences), in the style of Gul and Pesendorfer (2001) for example, would not alleviate the behavioral difference identified above, since the independence axiom in the menu framework that is satisfied by such models essentially implies the condition given in footnote 22. Moreover, in temptation models, the decision maker may decide even if he were paid to defer (ie. if the cost of deferral were negative), whereas this is not a consequence of the confidence-based approach.

24. The posteriors p_i could also be centres of centered confidence rankings; however, since ex post the decision maker is forced to choose, he will act like DM2 above, so the rest of the confidence rankings are irrelevant. Likewise, the distribution q over I may be thought of as a centre of a confidence ranking, but we may ignore the rest of the confidence ranking for the purposes of this simple exercise.

The first term on the first line is the expected value of the choice made by the person deferred to, which incorporates the cost of deferring x. The second term is the value of deciding rather than deferring for DM3: it is the standard expected value of the best act, but, unlike the equivalent term in (VIII.3), it incorporates the psychological cost of deciding from A when one does not hold the relevant beliefs with sufficient confidence for the decision, $c(D(A))$. (Recall that this cost of deciding is in utility units and is normalized to be non-negative.) The second line follows from the first by the definition of VI_x (equation (VIII.3)). So DM3 chooses according to a representation that is the same as (VIII.2) except that $c(D(A)) \leqslant u(x)$ is replaced by $VI_x(A) + c(D(A)) \leqslant 0$. Axiomatisation of this model and generalizations beyond the terms of this simple analysis are left to future research.

The incorporation of confidence into a standard information-based model of deferral thus introduces a second term – the cost of deciding from representation (VIII.2) – into the determinant of whether the decision maker defers. Unlike the value of information term, the cost of deciding term does not depend on the information differential between the decision maker and the person who will take the decision if deferred, but only on the decision itself and the level of confidence appropriate for it. It increases the value of deferral, and may thus lead to deferral in situations where the information differential alone would not.

This effect can be clearly seen on the demand value of deferral (that is, the highest amount the decision maker is willing to pay to defer). A decision maker incorporating both information- and confidence-related considerations into his choices, and who is not sufficiently confident in his beliefs given the decision, has a demand value of deferral of y for the choice from a menu A, where:

(VIII.7) $$VD_y(A) = VI_y(A) + c(D(A)) = 0$$

Contrasting this with the demand value under the pure information-based model (given by (VIII.4)) reveals that the incorporation of confidence bumps up the amount the decision making is willing to pay to defer. Moreover, this increase is greater for more important decisions, which call for more confidence. [25]

25. Basic comparative statics of the demand value of deferral can be immediately read off from equation (VIII.7) and the analysis in Appendix VIII.A. Decision makers that are more decision averse (see Appendix VIII.A for a precise definition of this notion) have higher costs of deciding from a menu A ($c(D(A))$), so their demand value of deferral for A is higher.

The introduction of confidence into the standard information-based approach may thus be able to explain cases where the information differential between the decision maker and the person to whom he defers does not justify the cost of deferral, but where the decision maker defers nonetheless. For instance, active portfolio management appears to perform too poorly with respect to passive benchmarks to justify its cost (Malkiel, 1995; Gruber, 1996), making it difficult to rationalize delegation solely on the basis of the difference in information (or information-gathering ability) between the investor and the portfolio manager. In such cases, where the decision is important and hence the cost of deciding is high, this cost may drive up the price that the investor is willing to pay to defer, ultimately leading to the choice to delegate.

Such cases could be just instances of a more general sort of 'calibration error'. The analysis suggests that standard estimations of the information differential based on willingness to pay to delegate or defer tend to be overvalued: part of what may be driving deferral is the reluctance to choose in important decisions on the basis of beliefs in which one has insufficient confidence. Correct estimations of the information differential may need to take this factor into account.

Finally, we note that the introduction of confidence into the information-based approach might also have consequences for what delegation or deferral options the decision maker adopts when he defers. If investors in fact see themselves as buying two services from portfolio managers – information, but also the taking of a difficult decision on their behalf – the relative importance of the two aspects in the decision to delegate may impact upon the competitive pressure on managers to provide high quality information. Exploration of such possible consequences of the introduction of confidence is left for further research. [26]

26. Although the details will of course differ considerably, there is little reason to expect that the general conclusions are specific to information-based approaches. There may even be interesting consequences of integrating confidence into accounts of deferral, or procrastination, based on time inconsistency, which have a different structure and often different primitives. For example, for a single task – such as a single decision to be taken – with a fixed cost, procrastination tends to decrease with increases in the per-period benefit (O'Donoghue and Rabin, 2001), suggesting that there would be less deferral in important decisions. Adding confidence to the model would imply that the cost (which is exogenous in the cited model) increases with the importance of the decision, tempering the mitigating effect of increased importance on deferral. This may contribute to explaining procrastination in important decisions such as the choice of retirement policy.

VIII.5 Related literature

Several related papers have already been mentioned previously; we now discuss some points of related literature that have not been treated above.

As noted in the Introduction, Bewley's (2002) theory of status quo choice – which can be thought of as deferral to nature – is one of the first axiomatic theories of intrinsic deferral-like phenomena. The essential intuition, which is behind part of the subsequent axiomatic literature on the status quo bias (for example, Masatlioglu and Ok (2005); Ortoleva (2010); Riella and Teper (2014))[27], is that one sticks with the status quo in the absence of appropriate determinate preference. Many discussions of deferral in behavioral economics (for example, Tversky and Shafir (1992)) and marketing (for example, Dhar (1997)), though generally couched in the context of decision under certainty, make reference at times to apparently similar intuitions that deferral results from preferences being incompletely constructed or 'conflicting'. To the extent that it is the decision maker's actual preferences that drive deferral, these are involved with intrinsic deferral. Although the theory proposed in this paper may be interpreted as a theory of status quo choice, the interpretation is more general, covering and focussing on cases of deferral to one's later self, and deferral to someone else.

As suggested previously, there is a relationship between representation (VIII.2) and incomplete preferences. Rather than discussing the whole literature on incomplete preferences, we focus on several conceptually and formally related papers.[28] Danan (2003b) and Kopylov (2009) both involve incomplete preference relations where incompleteness is interpreted in terms of postponing a decision. The former paper cashes out this option in terms of preference for flexibility and hence, as discussed in Sections VIII.1.1 and VIII.4.1, is concerned with extrinsic deferral (see Danan (2003a) for a related study in a choice-theoretical framework). The latter paper, which uses the unanimity representation à la Bewley and is noncommittal about the relationship between postponing a decision and preference for flexibility, may be interpreted in terms of intrinsic deferral. However, it does not offer an account of behaviour when deferral or postponement is costly.

The closest paper in this literature is Hill (2014), which, though it does not involve a deferral interpretation of incompleteness, can technically be thought of as a special case of

27. A technically related paper drawing on a different motivation is Savochkin (2014).
28. Readers seeking a more developed discussion of this literature are referred to Hill (2014).

(VIII.2) where deferral is free and all menus are binary. However, like Hill (2013), it assumes an exogenously-given notion of the stakes involved in, or importance of, a decision. This assumption is objectionable, especially if one considers the importance of a decision to be a subjective judgement on the part of the decision maker. No such assumption is made here. Indeed, as discussed in Remarks VIII.1 and VIII.3, a notion of stakes can be defined from the endogenous elements of representation (VIII.2). Theorem VIII.1 can thus be thought of as providing a behavioural foundation for the notion of stakes assumed in the cited papers, and hence an answer to the aforementioned criticism of the confidence-based approach that they set out. Indeed, this could be considered a separate contribution of the theorem; it is, to our knowledge, the first behavioural foundation for the notion of stakes in the literature.

The cost of deciding involved in representation (VIII.2) is reminiscent of cost factors involved in representations studied by Ergin and Sarver (2010); Ortoleva (2013); Buturak and Evren (2014). In a menu-preference framework, Ergin and Sarver (2010) derive endogenously the elements of a costly information acquisition representation, where different information structures have costs, called 'contemplation costs'. Given the conceptual proximity to the literature on information acquisition and preference for flexibility, the general remarks concerning the relationship between the approach taken here and standard approaches (see Sections VIII.1.1 and VIII.4.1) continue to apply. In particular, when applied to the question of deferral, Ergin and Sarver (2010) evidently take an extrinsic approach, which, even ignoring technical differences and differences in framework, clearly yields a different representation from that proposed here.

Buturak and Evren (2014), working in a choice-theoretic framework with a fixed 'default option' that is always present and no explicit cost of taking that option, introduce several related models incorporating a value of this option. The default option in their models can be either interpreted as a (fixed) status quo or as the opportunity to postpone the decision. In the most structured 'subjective state space' representation, both interpretations involve consideration of the decision maker's ex ante uncertainty about his future tastes. Hence, when interpreted in terms of deferral, this representation adopts an extrinsic approach, as opposed to the intrinsic approach taken here. The less structured 'variable threshold' representation, under the status quo interpretation of the default option, is related to the status quo literature mentioned above, and many of the points made there apply to this case. Behavioural differences between these representations and our representation (VIII.2) include the fact that the former demand that the decision maker defers from some singleton menus, whereas the latter implies that the decision maker never defers from singleton menus. Moreover, once this aspect is set aside, the independence axiom involved in the former representations appears to be stronger than our VIII-A3, to the extent that it demands that the decision maker choose from a mixture of a menu whenever he chooses from the menu itself (see also the discussion in Section VIII.3.2.1). [29]

Ortoleva (2013) proposes a representation of preferences over menus where the valuation of each menu differs from the standard 'additive EU representation' (Kreps, 1979; Dekel et al., 2001) by the addition of a 'thinking cost' that is a function of the menu. Despite the difference in interpretation – the thinking cost is generally motivated in terms of computational constraints, rather than the importance of the decision – the cost of deciding involved in representation (VIII.2) could be thought of as a possible source of Ortoleva's thinking cost, insofar as it represents the psychological cost of taking certain decisions. However, even putting aside technical differences and differences in framework, the representation in Ortoleva (2013) is quite different from ours, because it is concerned with a different choice problem. Ortoleva considers the problem of deciding which choice the decision maker would like to face at some future moment – and the cost of deciding or thinking in the choice he eventually faces will be a factor in this ex ante decision, though it does not affect the ex post choice taken. By contrast, we consider behavior in a single choice situation where deferral is an option – and the cost of deciding may influence the decision to defer or not. [30]

29. More precisely, in their models, the decision maker chooses from the mixture of a menu with an alternative if he chooses from the menu and from the singleton set containing the alternative; as noticed, this latter condition is automatically satisfied in the model proposed here. Another way of seeing the difference is by noting that, though our representation (VIII.2) can be written in the same form as their VT representation (their (1)), the threshold function will not be affine, as assumed in their representation (Definition 1).

30. This difference can be seen in the representations themselves, by considering the choice between f and g in the presence of a free deferral option: written in menu language, this is the choice among $\{f\}$, $\{g\}$ and $\{f, g\}$. Under Ortoleva's representation, there is no cost of thinking involved in the valuation of $\{f\}$ and $\{g\}$, but there may be a positive cost attached to $\{f, g\}$; under our representation, opting for $\{f\}$ or $\{g\}$ incurs a cost of deciding, whereas deferral – $\{f, g\}$ – incurs no cost. This is due to the difference in the choice problems considered. Since Ortoleva is considering what choice the decision maker would like to be faced with, if he is faced with $\{f\}$, no thinking will be required, so the cost is zero. Since we are considering behaviour in the choice between f and g with the possibility of deferring, $\{f\}$ amounts to making a decision

VIII.6 Conclusion

One potential source of deferral is lack of confidence in the beliefs needed to decide. Whilst recognised in some studies of deferral behavior – such as deferral to nature or to the status quo – this source is completely absent from many accounts of deferral – in particular standard accounts of delegation to others or deferral to one's future self. This paper makes a start at incorporating it into our models of such cases of deferral.

We develop a theory of deferral in which the decision maker's confidence in his beliefs plays a double role. It drives deferral, insofar as the decision maker is inclined to defer only when he lacks sufficient confidence in the relevant beliefs given the decision to be taken; whenever he has sufficient confidence, he decides. It prices deferral, to the extent that the value of the deferral option is determined by the confidence level appropriate for the decision under consideration. The theory applies naturally to situations where deferral is costly, which are the norm in cases of deferral to others or to one's future self. In these situations, a decision maker represented by the model decides if he has sufficient confidence in the relevant beliefs given the decision; if he lacks sufficient confidence, he decides whenever the cost of deferral outweighs its value, as determined by the confidence level appropriate for the decision, and he defers if not.

A behavioural axiomatization of this model is given, in which all the elements are derived endogenously. Choice-theoretically, the theory is quite reasonable: the essential axiomatic difference from the standard Savagean theory of decision under uncertainty (applied in a choice-theoretic setting) is to weaken the equivalent of the preference-theoretic completeness, transitivity and independence axioms to allow the decision maker to defer in cases where the standard axioms would have demanded decision.

The formulation of the model and its axiomatic analysis clarifies the relationship between confidence as a source of deferral and standard mechanisms for deferral to others or to one's future self considered in the literature, such as those based on expected difference or asymmetry of information between the deferrer and the deferree. In particular, it reveals behavioural patterns exhibited by information-based models but not by the proposed confidence-based one, hence showing that the two approaches to deferral are behaviourally distinct.

---

and hence may incur a cost.

Finally, the proposed model provides a starting point for the incorporation of confidence into standard approaches to deferral. A straightforward integration of confidence into an information-differential model yields a simple additive formula for the effect of confidence on the value of deferral. The model suggests that the introduction of confidence as a source of deferral increases the price the decision maker is willing to pay to defer, and that this effect is stronger for more important decisions (which call for more confidence). This may help explain some cases that are problematic for the standard approach: although the expected information differential may be insufficient to explain the observed willingness to pay for deferral or delegation, lack of confidence in the relevant beliefs, combined with the importance of the decision, would imply deferral even at relatively high costs.

VIII.A Attitudes to decision and deferral

In this Appendix, we perform some basic comparative statics on representation (VIII.2), bringing out the role of the various elements in the decision maker's attitude to deciding or deferring.

In the current context, there is a particularly simple comparison of decision makers' attitude to deciding: if decision maker 2 decides, rather than deferring, in every choice situation in which decision maker 1 decides, then decision maker 2 has less trouble deciding, and seeks deferral less, than decision maker 1. We shall say that decision maker 2 is less *decision averse*, or equivalently less *deferral seeking*, than decision maker 1. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *more decision averse* than γ^2 when, for all $A \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, if $\dagger_x \notin \gamma^1(A, \dagger_x)$ then $\dagger_x \notin \gamma^2(A, \dagger_x)$.

Beyond this benchmark comparison, we consider two others, which shall be useful in the separation of the roles of the different elements of the model. The first involves the comparison of when the cost of deferral is motivating. Recall from the discussion in Section VIII.3.2 that a cost is said to be motivating for a menu when it can drive the decision from it: there is a version of the menu from which the decision maker decides when deferral has this cost, though he does not decide when deferral is free. We shall say that decision maker 2 is more *cost motivated* than decision maker 1 if whenever a cost is motivating for a menu for decision maker 1, it is motivating for decision maker 2. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *less cost motivated* than γ^2 when, for all $A \in \wp(\mathcal{A})$ and $x \in \Re_{\geqslant 0}$, if $\overline{\gamma^1}^x(A) = \varnothing$ then $\overline{\gamma^2}^x(A) = \varnothing$.[31]

Decision aversion compares, for a given menu, the costs of deferral at which the decision makers decide. Another possible comparison focusses on the menus for which the cost is motivating, rather than the cost itself. If a decision maker decides from a menu A at every cost which is motivating for some other menu B, then this provides an indication about his deferral behaviour, relative to the extent to which cost motivates his choices. The more menus B for which this holds, the more he is inclined to decide rather than defer from A. If whenever decision maker 1 chooses from a menu at every cost that is motivating for another menu, decision maker 2 does the same, then there is a sense in which

31. Recall from the discussion in Section VIII.3.2 that, for a choice correspondence for costly deferral γ, $\overline{\gamma}^x(A) = \varnothing$ exactly when x is motivating for A.

decision maker 2 exhibits less decision aversion. We shall say that decision maker 2 is less *motivation-calibrated decision averse* than decision maker 1; the terminology reflects the fact that decision maker's cost motivation (rather than the cost itself) is used as a yardstick in the comparison. Formally, we say that (a decision maker whose choice correspondence is) γ^1 is *more motivation-calibrated decision averse* than γ^2 when, for all $A, B \in \mathcal{A}$, if $\dagger_x \notin \gamma_1(A, \dagger_x)$ for all $x \in \Re_{\geqslant 0}$ such that $\overline{\gamma^1}^x(B) = \varnothing$, then $\dagger_x \notin \gamma_2(A, \dagger_x)$ for all $x \in \Re_{\geqslant 0}$ such that $\overline{\gamma^2}^x(B) = \varnothing$.

The weakness of these notions – none of them ask that the decision makers choose the same acts when they decide – mean that the implications for the relationships between the elements of the representation are also weak. For a finer analysis, we shall follow a standard strategy in the literature and make assumptions sufficiently strong to guarantee that the decision makers have the same utilities and beliefs. To this end, we introduce the following notion.

Definition VIII.3. Let γ satisfy axioms VIII-A1–VIII-A10. The *confidence-in-choice* relation $\preccurlyeq$ on $\mathcal{A} \times \wp(\mathcal{A})$ is defined as follows: for any $A, B \in \wp(\mathcal{A})$ and $f, g \in \mathcal{A}$, $(f, A) \preccurlyeq (g, B)$ iff, for all $x \in \Re_{\geqslant 0}$, if $f \in \overline{\gamma}^x(A)$, then $g \in \overline{\gamma}^x(B)$.

Recall from Section VIII.3.2 that a cost is not motivating for a menu if, whenever the decision maker decides at this cost, he would decide even if deferral was free: he is, so to speak, sure enough to choose even if there were no cost to deferral. So if a cost x is not motivating for the choice from B and g is chosen, but x is motivating for the choice from A, then this can be taken as an indication that the decision maker is *more confident in his choice of g from B than in his choice from A*. Definition VIII.3 introduces a confidence-in-choice relation that, in these situations, ranks the choice of g from B ((g, B)) higher than any choice from A (for example, (f, A)).

We shall say that two decision makers are confidence equivalent if they have the same confidence-in-choice relation.

Definition VIII.4. Let γ^1 and γ^2 satisfy axioms VIII-A1–VIII-A10. γ^1 and γ^2 are *confidence equivalent* if $\preccurlyeq^1 = \preccurlyeq^2$.

Proposition VIII.A.1. *Let γ^1 and γ^2 satisfy axioms VIII-A1–VIII-A10, and be represented by (u_1, Ξ_1, D_1, c_1) and (u_2, Ξ_2, D_2, c_2) respectively. γ^1 and γ^2 are confidence equivalent if and only if u_2 is a positive affine transformation of u_1 and $\Xi_1 = \Xi_2$.*

This proposition confirms that a decision maker's confidence in his choices is entirely determined by his utilities and his confidence in beliefs. This is to be expected: to the extent that his choices are dictated by his utilities and his beliefs, it is reasonable that confidence in choices be determined by confidence in utilities – which is trivial in this model, because of the use of a single utility function – and confidence in beliefs, represented by the confidence ranking. Once differences in utilities and confidence in beliefs are accounted for, by comparing decision makers who have the same confidence in choices, the different comparisons mentioned above are entirely characterised by the relationship between the cautiousness coefficients and the cost functions, as the following proposition shows.

Proposition VIII.A.2. *Let γ^1 and γ^2 satisfy axioms VIII-A1–VIII-A10 and be confidence equivalent. Then:*

(i) *γ^1 is more decision averse than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $c_2(D_2(A)) \leqslant c_1(D_1(A))$ for all $A \in \overline{\Upsilon_2}$ and $D_2(A) \subseteq D_1(A)$ for all $A \notin \overline{\Upsilon}_2$.*

(ii) *γ^1 is less cost motivated than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $c_2(\mathcal{C}) \leqslant c_1(\mathcal{C})$ for all $\mathcal{C} \in \Xi$.*

(iii) *γ^1 is more motivation-calibrated decision averse than γ^2 if and only if there exist representations (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) of γ^1 and γ^2 respectively such that $D_2(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$.*

Since the cost function is unique, given the utility function, the comparison in part (ii) holds for all representations using the same utility function. Similarly, since the cautiousness coefficient D is only unique on $\overline{\Upsilon}$, it really only makes sense to compare D_1 and D_2 on $\overline{\Upsilon_1} \cap \overline{\Upsilon_2}$, and the comparisons on these sets (in parts (i) and (iii)) hold for all representations using the same utility function. (Note that both notions of decision aversion in fact imply that $\overline{\Upsilon_1} \supseteq \overline{\Upsilon_2}$, as is clear from the proof.) The formulation used for part (iii) in particular emphasises that the restriction to $\overline{\Upsilon_2}$ in the uniqueness of the cautiousness coefficient does not complicate applications: insofar as the behavioural consequences of comparisons of motivation-calibrated decision aversion are concerned, one can assume that the cautiousness coefficient containment condition holds everywhere.

Proposition VIII.A.2 brings out both the characterisation of the standard notion of decision aversion and the behavioural effects of the cautiousness coefficient and the cost function. More decision averse decision makers assign a higher cost to (the confidence level associated with) every menu. For each menu, they value the option of deferring from that menu more than less decision averse counterparts. Whilst not surprising, this does not separate the role of the cautiousness coefficient and the cost function, insofar as the cost assigned to a menu depends both on the level of confidence deemed appropriate for choice from the menu (determined by the cautiousness coefficient) and the cost of deciding when this level of confidence is required (determined by the cost function).

The other two parts of the result provide this separation. Comparison in terms of the cost function alone corresponds precisely to a difference in cost motivation: decision makers who are less cost motivated have a higher cost of deciding for every confidence level. It is the degree to which certain costs motivate decision makers' decisions that reveals the relationship between their costs of deciding at given confidence levels. By contrast, ordering in terms of motivation-calibrated decision aversion is characterised precisely in terms of an appropriate ordering of the cautiousness coefficients. It is intuitive that decision makers who assign higher levels of confidence to each menu are more decision averse: part (iii) of Proposition VIII.A.2 shows that this is the case, when decision aversion is measured on the motivation rather than the cost yardstick. Note that these results confirm the interpretations adopted for the elements of the model. Confidence drives deferral – and the cautiousness coefficient reflects tastes for choosing on the basis of limited confidence, when measured on a scale that is cost independent (the scale provided by motivation). Confidence also prices deferral – and the cost function captures the value of deferring, or equivalently the cost of deciding, fleshed out entirely in terms of the sensitivity of deferral behaviour to the cost of deferral.

Naturally, piecing together parts (ii) and (iii), one obtains a behavioural characterisation of when two decision makers satisfy both the containment condition on cautiousness coefficients and the ordering condition on cost functions: one is both less cost motivated and more motivation-calibrated decision averse than the other. It is evident from the proposition that these two comparisons imply that the decision maker is more decision averse. As Proposition VIII.C.3 in Appendix VIII.C shows, apart from this implication, the comparisons are independent: no other pair of comparisons imply the third.

VIII.B Proof of Theorem VIII.1

Throughout the remaining Appendices, B will denote the space of real-valued functions on S, and $ba(S)$ will denote the set of additive real-valued set functions on S, both under the Euclidean topology. B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $x \in \Re$, let x^* be the constant function taking value x.

The main part of Theorem VIII.1 consists in showing the sufficiency of the axioms ((i) to (ii) direction), the proof of which proceeds as follows. We first construct a binary relation $\leqq$ on $\wp(\mathcal{A})$ satisfying appropriate properties. This essentially orders menus according to whether a higher or lower confidence level is appropriate. The essential idea of the construction is that, for $A, B \in \Upsilon$, $A \leqq B$ iff $\iota(A) \leqslant \iota(B)$; the main work in the construction (Lemma VIII.B.1) is to extend this definition to the whole of $\wp(\mathcal{A})$ in such a way as the retain the appropriate properties. Then we establish (Lemma VIII.B.2) the representation for the case where deferral is free. To this end, for each indifference class r under $\leqq$ – which, recall, can be thought of as corresponding to a level of confidence – we define a function γ_r from menus to (perhaps empty) subsets, which can be thought of as representing the choices from menus considered 'as if' the appropriate level of confidence was r. We show (Lemmas VIII.B.4 to VIII.B.6) that for every non-minimal r there is a closed convex set of probability measures $\mathcal{C}_r$ representing γ_r in the sense that the elements selected by γ_r are the optimal ones according to the unanimity rule applied with $\mathcal{C}_r$. Lemmas VIII.B.8 and VIII.B.9 show that the choice correspondence for minimal r can be represented according to the unanimity rule with the intersection of the $\mathcal{C}_r$ for the other stakes levels, and that this set is a singleton. By Lemma VIII.B.7, the $\mathcal{C}_r$ form a nested family of sets, and we thus have a confidence ranking. By Lemmas VIII.C.8 and VIII.C.9, this confidence ranking is continuous, and by Lemma VIII.C.10, it is strict. By construction, the function that assigns to any level r the set $\mathcal{C}_r$ is a well-defined cautiousness coefficient. Moreover, the function that assigns to any set $\mathcal{C}_r$ in the confidence ranking the utility of $\iota(A)$ for any $A \in r$ is a well-defined cost function. Finally we show that these elements correctly capture the part of the representation where deferral has non-zero cost (Lemma VIII.B.10). We detail the main steps below; proofs of the technical lemmas are relegated to Appendix VIII.C.

VIII.B.1 Sufficiency of axioms

Note firstly that since $X = \Re$ is a separable metric space, the space of all Borel probability measures on X, and hence $\Delta(X)$ (which, recall from Section VIII.2.1, is the subspace of finitely additive probability measures on X), can be equipped with a separable metric (Billingsley, 2009, p72; Aliprantis and Border, 2007, Theorem 15.12). Hence, $\mathcal{A}$ is a separable metric space, equipped with the product metric, and $\wp(\mathcal{A})$ is a metric space, under the Hausdorff metric. If $\Phi = \varnothing$, then by VIII-A4, $\dagger_x \notin \gamma(A, \dagger_x)$ for all $A \in \wp(\mathcal{A})$, $x \in \Re_{\geqslant 0}$, and the result follows from the standard representation theorem for expected utility. Throughout the rest of the proof, we thus assume that $\Phi \neq \varnothing$. We begin by stating some preliminary results (see Appendix VIII.C for proofs).

The first is a generalisation of Hill (2009, Theorem 2), which held for finite menus, to the case of infinite menus (with topological structure).

Theorem VIII.2. *Let X be a metric space and $\wp(X)$ the set of non-empty compact subsets of X, endowed with the Hausdorff topology. Let $\gamma : \wp(X) \to 2^X$ be a function such that $\gamma(A) \subseteq A$ for all $A \in \wp(X)$. Suppose that γ satisfies the following continuity property: for all sequences $A_i \in \wp(X)$ and $x_i \in X$ with $A_i \to A$ and $x_i \to x$, if $x_i \in \gamma(A_i)$ for all i, then $x \in \gamma(A)$.*[32] *The following are equivalent:*

(i) There exists a reflexive, transitive binary relation $\preceq$ such that $\gamma(A) = \{x \in A \mid x \succeq y \ \forall y \in A\}$;

(ii) γ satisfies the following properties:

α	*if $x \in A \subseteq B$ and $x \in \gamma(B)$, then $x \in \gamma(A)$*
π	*if $x \in A$, $y \in A \cap B$, $y \in \gamma(B)$ and $x \in \gamma(A)$, then $x \in \gamma(A \cup B)$*
$sing$	*if $A = \{x\}$, then $x \in \gamma(A)$.*

Moreover, $\preceq$ is unique.

Finally, if γ always takes non-empty values, (ii) is equivalent to the existence of a reflexive, transitive and complete preference relation representing γ, even in the absence of the continuity assumption.

32. Note that γ is not necessarily a correspondence, because it is not necessarily non-empty-valued. The continuity condition, which is essentially a version of the standard upper hemi-continuity condition for correspondences, is thus stated explicitly.

Secondly, Lemma VIII.C.1 establishes that there is a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \to \Re$ representing the restriction of γ to menus of constant acts. We fix such a u for use throughout the proof, and let $K = u(\Delta(X))$ and $B(K)$ be the set of functions in B taking values in K. Note that since u is strictly increasing, K is an open interval. $B(K)$ is naturally isomorphic to a subset of $\Re^{|S|}$, and we take it to be equipped with the Euclidean topology. Lemma VIII.C.2 guarantees that, under u, γ generates a well-defined choice correspondence for costly deferral on $\wp(B(K)) \times \mathcal{D}$, which, with slight abuse of notation, we denote by γ.

VIII.B.1.1 Construction of 'higher level of confidence appropriate' relation

Lemma VIII.B.1. *There exists a non-trivial, continuous weak order*[33] $\leqq$ *on* $\wp(\mathcal{A})$ *satisfying:*

i. for all $A, B \in \Upsilon$, $A \leqq B$ *iff* $\iota(A) \leqslant \iota(B)$*;*

ii. for all $A \in \Upsilon$, $B \in \Phi \backslash \Upsilon$, *if* $\iota(A) > mmc(B)$, *then* $A > B$*;*

iii. for all $A, B, \in \wp(\mathcal{A})$, *if* $A \overset{e.e.}{\simeq} B$, *then* $A \equiv B$*;*

iv. for all $A, B \in \Phi \subseteq \wp(\mathcal{A})$, *there exists* $h, h' \in \mathcal{A}$ *and* $\alpha, \alpha' \in (0, 1]$ *such that* $A_\alpha h \leqq B \leqq A_{\alpha'} h'$*;*

v. there exists $A \in \Phi$ *such that, for all* $B \in \wp(\mathcal{A})$, $A \leqq B$.

where $\equiv$ *and* $<$ *are the symmetric and asymmetric parts of* $\leqq$ *respectively.*

Proof. Recall that, by assumption, Φ is non-empty. Note that, by VIII-A7, $\wp(\Delta(X)) \subseteq \Phi^c$. We first define the function $\tilde{\iota} : \Phi \to \Re$ as follows: $\tilde{\iota}(A) = \iota(A)$ if $A \in \Upsilon$ and $\tilde{\iota}(A) = mmc(A)$ otherwise. By VIII-A9 part iii. $\tilde{\iota}$ is continuous.

By Lemma VIII.C.1, u is a strictly increasing zeroed continuous affine utility function representing the restriction of γ to sets of constant acts. Note that $\{d \in \Delta(X) \mid \gamma(\{c, d\}, \dagger_0) = \{c\}\}$ and $\{d \in \Delta(X) \mid \gamma(\{c, d\}, \dagger_0) = \{d\}\}$ is a subbase for the set of open subsets of $\Delta(X)$ that are pre-images of sets in K under u; since the image under u of these sets is $\{y \in K \mid y < u(c)\}$ and $\{y \in K \mid y > u(c)\}$, a subbase for the set of open subsets of K, u is a quotient map (Hart et al., 2004, Section b-4). Moreover, there is a many-to-one map $\hat{u} : \mathcal{A} \to B(K)$, given by $\hat{u}(f)(s) = u \circ f(s)$, for all $f \in \mathcal{A}$, $s \in S$.

33. A weak order is a reflexive, transitive and complete binary relation; a relation $\leqslant$ on a topological space X is continuous if $\{y \in X \mid x \leqslant y\}$ and $\{y \in X \mid x \geqslant y\}$ are closed in X.

Since u is a quotient map and $\hat{u}$ is a finite product of quotient maps, $\hat{u}$ is also a quotient map. Let $\tilde{u} : \wp(\mathcal{A}) \to \wp(B(K))$ be the generated map between the compact subsets of $\mathcal{A}$ and $B(K)$ (since $\hat{u}$ is continuous, the image of each element of $\wp(\mathcal{A})$ is compact). By Aliprantis and Border (2007, Theorem 3.91), the Hausdorff topology on $\wp(\mathcal{A})$ coincides with the Vietoris topology on $\wp(\mathcal{A})$, and similarly for $\wp(B(K))$. The collection of sets of the form $G_0^h \cap G_1^l \cap \cdots \cap G_n^l$, where the G_i are open subsets of $B(K)$, $G^h = \{F \in \wp(B(K)) \mid F \subset G\}$ and $G^l = \{F \in \wp(B(K)) \mid F \cap G \neq \varnothing\}$, is a base for the Vietoris topology on $\wp(B(K))$ (Aliprantis and Border, 2007, Definition 3.89). So $H \subseteq \wp(B(K))$ is open iff it is the union of such sets, and this holds iff $\tilde{u}^{-1}(H)$ is the union of sets of the form $\tilde{u}^{-1}(G_0^h) \cap \tilde{u}^{-1}(G_1^l) \cap \cdots \cap \tilde{u}^{-1}(G_n^l)$ with G_i open in $B(K)$. But since $\tilde{u}^{-1}(G^h) = \{\tilde{u}^{-1}(F) \mid F \in \wp(B(K)),\ F \subset G\} = \bigcup_{G' \in \hat{u}^{-1}(G)} \{F \in \wp(\mathcal{A}) \mid F \subset G'\} = \bigcup_{G' \in \hat{u}^{-1}(G)} (G')^h$ and $\tilde{u}^{-1}(G^l) = \{\tilde{u}^{-1}(F) \mid F \in \wp(B(K)),\ F \cap G \neq \varnothing\} = \bigcup_{G' \in \hat{u}^{-1}(G)} \{F \in \wp(\mathcal{A}) \mid F \cap G' \neq \varnothing\} = \bigcup_{G' \in \hat{u}^{-1}(G)} (G')^l$, $\tilde{u}^{-1}(H)$ is the union of sets of the form $\tilde{u}^{-1}(G_0^h) \cap \tilde{u}^{-1}(G_1^l) \cap \cdots \cap \tilde{u}^{-1}(G_n^l)$ with G_i open in $B(K)$ iff it is a pre-image of an open subset of $\wp(B(K))$ under $\tilde{u}$ that is the union of sets of the form $(G_0')^h \cap (G_1')^l \cap \cdots \cap (G_n')^l$ where the G_i' are open subsets of $\mathcal{A}$. By the definition of the Vietoris topology on $\wp(\mathcal{A})$, this latter property holds iff $\tilde{u}^{-1}(H)$ is open. So H is open in $\wp(B(K))$ iff $\tilde{u}^{-1}(H)$ is open in $\wp(\mathcal{A})$. Hence $\tilde{u}$ is a quotient map between $\wp(\mathcal{A})$ and $\wp(B(K))$ taken with the Vietoris, or equivalently Hausdorff, topology.

Recall that, with slight abuse of notation, we use $\gamma : \wp(B(K)) \times \mathcal{D} \to 2^{B(K) \cup \mathcal{D}} \backslash \varnothing$ to denote the choice correspondence for costly deferral generated on $\wp(B(K)) \times \mathcal{D}$ by γ under $\tilde{u}$; similarly, we use ι to denote the image of ι and likewise for mmc, Υ, Φ and $\tilde{\iota}$. (Lemma VIII.C.2 and VIII-A6 imply that γ is well-defined on $B(K)$ and that ι, mmc, $\tilde{\iota}$, Υ and Φ are well-defined on $\wp(B(K))$.) Since $\tilde{u}$ is a quotient map, it follows from a standard result (Kelley, 1975, Theorem 3.9) that the continuity of $\tilde{\iota}$ implies that the function it generates on $\Phi \subseteq \wp(B(K))$ is continuous.

Let $\underline{x} = \inf \iota(\Upsilon)$; by the definition of $\tilde{\iota}$, $\inf \tilde{\iota}(\Upsilon) = \inf \tilde{\iota}(\Phi) = \underline{x}$. For each $A \in \Phi$ such that $mmc(A) \neq \underline{x}$, $\{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\} \cap (\overline{\Upsilon})^c \neq \varnothing$, by VIII-A10 part ii. For each such A, let $\widehat{A}$ be any element of $\{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\}$ such that $d(\hat{A}, \overline{\Upsilon}) > \frac{1}{2} \sup_{A' \in \{\alpha A + (1-\alpha) l \mid \alpha \in (0,1], l \in B(K)\}} d(A', \overline{\Upsilon})$ (where d is the Euclidean metric on $B(K)$). Let $\mathcal{B} = \overline{\{\widehat{A} \mid A \in \wp(B(K)),\ mmc(A) \neq \underline{x}\}}$. Note that, by VIII-A9 part i., $A \notin \Upsilon$ for all $A \in \mathcal{B}$. We now show that, if $\mathcal{B} \cap \overline{\Upsilon} \neq \varnothing$, then $\tilde{\iota}(\mathcal{B} \cap \overline{\Upsilon}) = \underline{x}$. For any $A \in \mathcal{B} \cap \overline{\Upsilon}$, the definition of $\mathcal{B}$ and the continuity of the metric implies that there is no $\alpha A + (1-\alpha)l$, for $\alpha \in (0,1]$, $l \in B(K)$, with distance from $\overline{\Upsilon}$ strictly greater than the distance of A from $\overline{\Upsilon}$, whence, by VIII-A10 part ii., $mmc(A) = \underline{x}$. Since $A \notin \Upsilon$, it follows from the definition of $\tilde{\iota}$ that $\tilde{\iota}(A) = mmc(A) = \underline{x}$, as required.

Now consider the correspondence $t : \Phi \rightrightarrows \Re$ (ie. map from $\Phi \to 2^{\Re} \backslash \varnothing$) defined as follows: $t(A) = \{\tilde{\iota}(A)\}$ if $A \in \overline{\Upsilon}$; $t(A) = \{\underline{x}\}$ if $A \in \mathcal{B}$; and $t(A) = [\underline{x}, \tilde{\iota}(A)]$ otherwise. t is obviously a well-defined closed and convex-valued correspondence. We now show that it is lower hemicontinuous. By Aliprantis and Border (2007, Theorem 17.21), it suffices to show that for any $(A_n) \in \Phi$ with $A_n \to A$, and any $y \in t(A)$, there exists $y_n \to y$ with $y_n \in t(A_n)$ for all n. Whenever $A \notin \mathcal{B} \cap \overline{\Upsilon}$, this is an immediate consequence of the continuity of $\tilde{\iota}$, the fact that $\overline{\Upsilon}$ and $\mathcal{B}$ are closed, and of the definition of t. Whenever $A \in \mathcal{B} \cap \overline{\Upsilon}$, as shown above, $\tilde{\iota}(A) = \underline{x}$, so the continuity of t at A follows from the continuity of $\tilde{\iota}$. Hence t is lower hemicontinuous.

Since Φ is metric, and hence paracompact, the Michael Selection Theorem (Aliprantis and Border, 2007, Theorem 17.66) implies that there exists a continuous function $s : \Phi \to \Re$ such that $s(A) \in t(A)$ for all $A \in \Phi$. By definition s agrees with $\tilde{\iota}$ on $\overline{\Upsilon}$, and it attains its infimal value.

It remains to extend s to $B(K)$. Since $\wp(B(K))$ is a metric space, so is Φ. Since s is continuous on this space, by Aliprantis and Border (2007, Lemma 3.12) there exists an equivalent metric on Φ under which s is Lipschitz, and hence uniformly continuous. By Aliprantis and Border (2007, Lemma 3.11), there is a uniformly continuous extension of s to the closure of Φ under this metric, which, since the metrics are equivalent, coincides with $\overline{\Phi}$, the closure under the Hausdorff metric. Let $\hat{s}$ be this extension; it is a continuous function (under the Hausdorff metric on $\wp(B(K))$) on $\overline{\Phi}$. Moreover, by definition, $\hat{s}(\overline{\Phi}) = \overline{s(\Phi)}$.

Since $\wp(B(K))$ is a metric space (and hence normal) and $\overline{\Phi}$ is closed, Tietze's extension theorem (Aliprantis and Border, 2007, Theorem 2.47) applies, and implies that there is a continuous real-valued function σ on $\wp(B(K))$ and agreeing with $\hat{s}$ on $\overline{\Phi}$. Moreover, σ can be chosen such that $\sigma(\wp(B(K))) \subseteq \overline{\hat{s}(\overline{\Phi})} = \overline{s(\Phi)}$.

Define the relation $\leqq$ on $\wp(\mathcal{A})$ as follows: for any $A, B \in \wp(\mathcal{A})$, $A \leqq B$ iff $\sigma(\tilde{u}(A)) \leqslant \sigma(\tilde{u}(B))$. Since σ is a continuous real-valued function defined everywhere on $\wp(B(K))$, $\leqq$ is a continuous weak order; since σ is defined on $\wp(B(K))$, $\leqq$ satisfies property iii. By definition, s and hence σ is not a constant function, so $\leqq$ is non-trivial. Moreover, s attains its infimal value, so σ attains its infimal value with an element in Φ, and hence $\leqq$ satsifies property v. We now show that it satisfies the other required properties.

Property i. follows immediately from the fact that, for all $A \in \wp(\mathcal{A})$ with $A \in \Upsilon$, by definition, $\sigma(\tilde{u}(A)) = \iota(A)$. As concerns property ii., if $A \in \Upsilon$ and $B \in \Phi \backslash \Upsilon$, then, by the definition of σ, $\sigma(\tilde{u}(A)) = \iota(A)$ and $\sigma(\tilde{u}(B)) \leqslant mmc(B)$. Hence if $\iota(A) > mmc(B)$, then $\sigma(\tilde{u}(A)) > \sigma(\tilde{u}(B))$, as required. We now show that $\leqq$ satisfies property iv. On the one hand, by VIII-A10 part i., for any $A, B \in \Phi$ with $\sigma(\tilde{u}(A)) < \sigma(\tilde{u}(B)) = x$, there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\sigma(\tilde{u}(A_{\alpha'}h')) = \iota(A_{\alpha'}h') \geqslant x$, as required. On the other hand, to treat the case where $A, B \in \Phi$ with $\sigma(\tilde{u}(A)) > \sigma(\tilde{u}(B)) = x$, we consider two cases. If $mmc(A) = \inf \iota(\Upsilon) = \underline{x}$, by VIII-A5 there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(\tilde{u}(A_{\alpha'}h'), \dagger_0)$, so, by the definition of σ, $\sigma(\tilde{u}(A_{\alpha'}h')) \leqslant mmc(A) = \underline{x} \leqslant x$. On the other hand, if $mmc(A) \neq \inf \iota(\Upsilon)$, by construction there exists $\alpha' \in (0,1]$ and $h' \in \mathcal{A}$ such that $\tilde{u}(A_{\alpha'}h') \in \mathcal{B}$ and so $\sigma(\tilde{u}(A_{\alpha'}h')) = \underline{x} \leqslant x$. It follows that $\leqq$ satisfies property iv, as required. This completes the proof of Lemma VIII.B.1.

□

VIII.B.1.2 Construction of confidence ranking and representation

Throughout this section, and the relevant Lemmas in Appendix VIII.C, $\leqq$ is assumed to be any relation satisfying the conditions i–v of Lemma VIII.B.1. We say that a function $D : \wp(\mathcal{A}) \to \Xi$ respects $\leqq$ if, for any $A, B \in \wp(\mathcal{A})$, $D(A) \subseteq D(B)$ iff $A \leqq B$. The following lemma is central.

Lemma VIII.B.2. *There exists a centred strict continuous confidence ranking Ξ and a cautiousness coefficient $D : \wp(\mathcal{A}) \to \Xi$ respecting $\leqq$ such that, for all $A \in \wp(\mathcal{A})$,*

$$\gamma(A, \dagger_0) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \{\dagger_0\} & \textit{if} \sup(A, u, D(A)) = \varnothing \end{cases} \tag{VIII.8}$$

where u is as in Lemma VIII.C.1.

Proof. By Lemma VIII.C.1, there exists a strictly increasing zeroed continuous affine utility function u representing choices on menus consisting of constant acts. We use the notation introduced in the proof of Lemma VIII.B.1; in particular, we denote by $\tilde{u} : \wp(\mathcal{A}) \to$

$\wp(B(K))$ the function on menus generated by u. As above, we use this function to map γ into a choice correspondence for costly deferral on $\wp(B(K)) \times \mathcal{D}$, which we also call γ. Let $\leqq$ on $\wp(B(K))$ be the image of $\leqq$ under u, and similarly for ι, mmc, Υ, Φ and $\bar{\gamma}^x$. We begin by stating two consequences of the construction and properties of $\leqq$.

Firstly, by Lemma VIII.C.4, for all $A \in \wp(B(K))$, $\alpha \in (0,1]$ and $l \in B(K)$, if $\alpha A + (1-\alpha)l \leqq A$ and $\dagger_0 \notin \gamma(A, \dagger_0)$, then $\dagger_0 \notin \gamma(\alpha A + (1-\alpha)l, \dagger_0)$, and similarly for the case where $\alpha A + (1-\alpha)l \geqq A$. Moreover, Lemma VIII.C.5 establishes that, for every $A, A' \in \Phi$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \equiv A'$.

Let $\mathcal{S}$ be the set of equivalence classes of $\leqq$. As standard, $\leqq$ on $\wp(B(K))$ generates a relation on $\mathcal{S}$, which will be denoted $\leqslant$ (with symmetric and asymmetric components $=$ and $<$ respectively): for $r, s \in \mathcal{S}$, $r \leqslant s$ iff, for any $A \in r$ and $A' \in s$, $A \leqq A'$. $r \in \mathcal{S}$ is a minimal (respectively maximal) element if $r \leqslant s$ (resp. $r \geqslant s$) for all $s \in \mathcal{S}$. Note that, since $\leqslant$ is a linear ordering, there is at most one minimal (resp. maximal) element. By property v of $\leqq$, there is a minimal element of $\mathcal{S}$, which we call $\underline{\mathcal{S}}$; the maximal element, if it exists, is denoted by $\overline{\mathcal{S}}$. We say that an element $r \in \mathcal{S}$ is *full* if, for every $A \in \Phi$, there exists $\alpha \in (0,1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \in r$. Let $\mathcal{S}^f$ be the set of full elements of $\mathcal{S}$. By Lemma VIII.C.5 and the construction of $\leqq$, every non-maximal element of $\mathcal{S}$ is full. Let $\mathcal{S}^+$ be the set of non-minimal elements of $\mathcal{S}^f$. For each $r \in \mathcal{S}^f$, define $\gamma_r : \wp(B(K)) \to 2^{B(K)}$ as follows: for all $A \in \Phi$, $\gamma_r(A) = B$ if there exists $B \in 2^{B(K)}$, $l \in B(K)$ and $\alpha \in (0,1]$ such that $\gamma(\alpha A + (1-\alpha)l, \dagger_0) = \alpha B + (1-\alpha)l$ and $\alpha A + (1-\alpha)l \in r$, and $\gamma_r(A) = \varnothing$ otherwise; and for all $A \notin \Phi$, $\gamma_r(A) = \gamma(A, \dagger_0)$. Lemma VIII.C.6 guarantees that $\gamma(\alpha A + (1-\alpha)l, \dagger_0) = \alpha B + (1-\alpha)l$ if and only if $\gamma(\beta A + (1-\beta)m, \dagger_0) = \beta B + (1-\beta)m$ for all $l, m \in B(K)$ and $\alpha, \beta \in (0,1]$ such that $\alpha A + (1-\alpha)l, \beta A + (1-\beta)m \in r$, so γ_r is well-defined for every $r \in \mathcal{S}^f$.

By Lemma VIII.C.7, the functions γ_r respect the ordering $\leqslant$ on $\mathcal{S}$ in the following sense: for all $A \in \wp(\mathcal{A})$ and $r, s \in \mathcal{S}^f$, if $r \geqslant s$, then $\gamma_r(A) \neq \varnothing$ implies that $\gamma_s(A) = \gamma_r(A)$. We now establish some further properties of the γ_r.

Lemma VIII.B.3. *For every $r \in \mathcal{S}^+$, γ_r satisfies the following continuity property: for all sequences $A_n \in \wp(B(K))$, $a_n \in B(K)$ with $A_n \to A$ and $a_n \to a$, if $a_n \in \gamma_r(A_n)$ for all $n \in \mathbb{N}$, then $a \in \gamma_r(A)$.*

Proof. Let r be a non-minimal element of $\mathcal{S}^f$, and let a_n, A_n be a pair of sequences with $A_n \to A$, $a_n \to a$ and $a_n \in \gamma_r(A_n)$ for all $n \in \mathbb{N}$. We consider the case where $A \in \Phi$, and suppose without loss of generality that $A \in r$. (The case where $A \notin \Phi$ can be treated in an analogous fashion, replacing A in the argument below by any member of r.) If there exists $N \in \mathbb{N}$ such that $A_n \leqq A$ for all $n \geqslant N$, then, by Lemma VIII.C.7, $a_n \in \gamma(A_n, \dagger_0)$ for all $n \geqslant N$, so by VIII-A9 part i., $a \in \gamma(A, \dagger_0)$, and thus $a \in \gamma_r(A)$, as required.

Now suppose that there is no such N. Since r is a non-minimal element of $\mathcal{S}^f$, there exists $l \in B(K)$ and $\bar{\delta} \in (0,1)$ such that $\bar{\delta}A + (1-\bar{\delta})l < A$. Let $\eta = \min\{\delta \in (\bar{\delta}, 1] | \ \delta A + (1-\delta)l \geqq A\}$ (by the continuity of $\leqq$ this is a minimum). Consider any $\delta \in (\bar{\delta}, \eta)$; by the definition of η, $\delta A + (1-\delta)l < A$. By the continuity of $\leqq$, there exists $N_\delta \in \mathbb{N}$ such that, for all $n > N_\delta$, $\delta A_n + (1-\delta)l < A$. Lemma VIII.C.7 implies that $\delta a_n + (1-\delta)l \in \gamma(\delta A_n + (1-\delta)l, \dagger_0)$ for every $n > N_\delta$, whence, by VIII-A9 part i., $\delta a + (1-\delta)l \in \gamma(\delta A + (1-\delta)l, \dagger_0)$. Since this holds for all $\delta \in (\bar{\delta}, \eta)$, VIII-A9 part i. implies that $\eta a + (1-\eta)l \in \gamma(\eta A + (1-\eta)l, \dagger_0)$. Since $\eta A + (1-\eta)l \in r$, it follows that $a \in \gamma_r(A)$, as required.

□

Lemma VIII.B.4. *For each $r \in \mathcal{S}^f$, there exists a unique reflexive, transitive binary relation $\preceq_r$ representing γ_r, in the following sense: $\gamma_r(A) = \{a \in A \mid a \succeq_r b \ \ \forall b \in A\}$. Moreover, if $r = \underline{\mathcal{S}}$, then $\preceq_r$ is complete.*

Proof. We first show that for every $r \in \mathcal{S}^f$, γ_r satisfies properties α and π in Theorem VIII.2 (see also Hill (2009)).

Consider firstly $a \in B(K)$ and $A, B \in \wp(B(K))$ such that $a \in A \subseteq B$ and $a \in \gamma_r(B)$, and let $x_r \in \Re_{\geqslant 0}$ be such that, for any $A' \in \wp(\mathcal{A})$ such that $A' \in \Upsilon$ and $A' \in r$, $\iota(A') = x_r$ (by property i of $\leqq$, x_r is well-defined). It follows from $a \in \gamma_r(B)$ that, for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$, $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$. To see this, take any $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$. If $\alpha B + (1-\alpha)l \in \Upsilon$, then $\iota(\alpha B + (1-\alpha)l) \leqslant r$, so $\alpha B + (1-\alpha)l \in s \leqslant r$, and since, by Lemma VIII.C.7, $\gamma_s(B) = \gamma_r(B)$, $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$, contradicting the assumption that $\alpha B + (1-\alpha)l \in \Upsilon$. So $\alpha B + (1-\alpha)l \notin \Upsilon$, whence it follows from $a \in \gamma_r(B)$ by the argument in Lemma VIII.C.6 that $\alpha a + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$, as required. So $a \in \bar{\gamma}^{x_r}(B)$, from which it follows, by VIII-A1, that $a \in \bar{\gamma}^{x_r}(A)$. If $A \notin \Phi$, then $\dagger_{x_r} \notin \gamma(A, \dagger_{x_r})$, so $a \in \bar{\gamma}^{x_r}(A)$ implies that $a \in \gamma(A, \dagger_0)$. Now consider the case where $A \in \Phi$, and suppose without loss of generality that $A \in r$. By the definition of ι, VIII-A4 and VIII-A9 part i., $\dagger_{x_r} \notin \gamma(A, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A)$, implies that $a \in \gamma(A, \dagger_0)$. So $a \in \gamma_r(A)$, and γ_r satisfies property α.

Now consider $a, b \in B(K)$ and $A, B \in \wp(B(K))$ such that $a \in A$, $b \in A \cap B$, $a \in \gamma_r(A)$ and $b \in \gamma_r(B)$. Let $x_r \in \Re_{\geqslant 0}$ be such that, for any $A' \in \wp(\mathcal{A})$ such that $A' \in \Upsilon$ and $A' \in r$, $\iota(A') = x_r$. As above, for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha A + (1-\alpha)l, \dagger_{x_r})$, $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$, and for all $\alpha \in (0,1]$, $l \in B(K)$ such that $\dagger_{x_r} \notin \gamma(\alpha B + (1-\alpha)l, \dagger_{x_r})$, $\alpha b + (1-\alpha)l \in \gamma(\alpha B + (1-\alpha)l, \dagger_0)$. So $a \in \bar{\gamma}^{x_r}(A)$ and $b \in \bar{\gamma}^{x_r}(B)$, whence, by VIII-A2, $a \in \bar{\gamma}^{x_r}(A \cup B)$. If $A \cup B \notin \Phi$, then $\dagger_{x_r} \notin \gamma(A \cup B, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A \cup B)$, implies $a \in \gamma(A \cup B, \dagger_0)$. Now consider the case where $A \cup B \in \Phi$ and suppose without loss of generality that $A \cup B \in r$. By the definition of ι, VIII-A4 and VIII-A9 part i., $\dagger_{x_r} \notin \gamma(A \cup B, \dagger_{x_r})$, which, given $a \in \bar{\gamma}^{x_r}(A \cup B)$, implies that $a \in \gamma(A \cup B, \dagger_0)$. So $a \in \gamma_r(A \cup B)$, and γ_r satisfies property π.

Now consider the case where r is non-minimal in $\mathcal{S}$. By Lemma VIII.B.3, γ_r satisfies the continuity condition in Theorem VIII.2. By VIII-A8 part ii., for every $a \in B(K)$, $a \in \gamma(\{a\}, \dagger_0)$, so $a \in \gamma_r(\{a\})$, and γ_r satisfies $sing$. Theorem VIII.2 implies that there exists a unique reflexive, transitive binary relation representing γ_r, as required.

For the case where $r = \underline{\mathcal{S}}$, note firstly that, by VIII-A5, for every $A \in \wp(\mathcal{A})$, there exists $\alpha \in (0,1]$, $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$. By Lemma VIII.C.7, it follows that $\gamma_{\underline{\mathcal{S}}}(A) \neq \varnothing$ for all $A \in \wp(\mathcal{A})$: $\gamma_{\underline{\mathcal{S}}}$ always takes non-empty values. Theorem VIII.2 implies the required representation.

□

Note that, by VIII-A8 part i., $\preceq_r$ is non-trivial for all $r \in \mathcal{S}$. We now establish some other properties of the relations $\preceq_r$. Recall that a binary relation $\preceq$ on $B(K)$ is

— *monotonic* if, for all $a, b, \in B(K)$, if $a \leqslant b$ then $a \preceq b$.

— *affine* if, for all $a, b, c \in B(K)$ and $\alpha \in (0,1)$, $a \preceq b$ iff $\alpha a + (1-\alpha)c \preceq \alpha b + (1-\alpha)c$.

— *continuous* if, for all $a_n, b_n \in B(K)$, if $a_n \preceq b_n$ for all $n \in \mathbb{N}$, $a_n \to a$ and $b_n \to b$, then $a \preceq b$.

Lemma VIII.B.5. *For every $r \in \mathcal{S}^f$, $\preceq_r$ is monotonic and affine. Moreover, if $r \in \mathcal{S}^+$, $\preceq_r$ is continuous.*

Proof. Monotonicity. Suppose that $a \leqslant b$. So $\alpha a + (1-\alpha)l \leqslant \alpha b + (1-\alpha)l$ for all $l \in B(K)$ and $\alpha \in (0,1]$. VIII-A8 part ii. implies that $\alpha b + (1-\alpha)l \in \gamma(\alpha\{a,b\} + (1-\alpha)l, \dagger_0)$ for all such l and α. It follows from the definition of γ_r and Lemma VIII.B.4 that $a \preceq_r b$.

Affineness. We consider the case where $\{a, b\}, \{\alpha a + (1-\alpha)c, \alpha b + (1-\alpha)c\} \in \Phi$; the other case is treated similarly. Note that it follows from the specification of the case and VIII-A8 part ii. that $a \neq b$. Since $r \in \mathcal{S}^f$, there exists $\beta \in (0, 1]$ and $l \in B(K)$ such that $\beta\{a, b\} + (1-\beta)l \in r$. Consider $\{\beta(\alpha a + (1-\alpha)c) + (1-\beta)l, \beta(\alpha b + (1-\alpha)c) + (1-\beta)l\}$: since $r \in \mathcal{S}^f$, there exists $\delta \in (0, 1]$ and $m \in B(K)$ such that $\delta\{\beta(\alpha a + (1-\alpha)c + (1-\beta)l, \beta(\alpha b + (1-\alpha)c + (1-\beta)l\} + (1-\delta)m \in r$. Note that $\delta(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m = \alpha\delta(\beta a + (1-\beta)l) + (1-\alpha\delta)(\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c + (1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m)$, where $\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c + (1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m \in B(K)$ since it is a mix of elements of $B(K)$; similarly for b. Let $f, g, h \in \mathcal{A}$ be such that $\beta a + (1-\beta)l = u \circ f$, $\beta b + (1-\beta)l = u \circ g$ and $\frac{\delta-\alpha\delta}{1-\alpha\delta}(\beta c + (1-\beta)l) + \frac{1-\delta}{1-\alpha\delta}m = u \circ h$. Since $\{f, g\} \equiv \{f_{\alpha\delta}h, g_{\alpha\delta}h\}$, by Lemma VIII.C.4 and VIII-A3, $\beta b + (1-\beta)l \in \gamma(\beta\{a, b\} + (1-\beta)l, \dagger_0)$ iff $\delta(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m \in \gamma(\{\delta(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m, \delta(\beta(\alpha b + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m\}, \dagger_0)$. But since $\delta(\beta(\alpha a + (1-\alpha)c) + (1-\beta)l) + (1-\delta)m = \beta\delta(\alpha a + (1-\alpha)c) + (1-\beta\delta)(\frac{\delta-\beta\delta}{1-\beta\delta}l + \frac{1-\delta}{1-\beta\delta}m)$, and similarly for b, it follows that $a \preceq_r b$ iff $\alpha a + (1-\alpha)c \preceq_r \alpha b + (1-\alpha)c$, as required.

Continuity. Continuity of γ_r for $r \in \mathcal{S}^+$ is an immediate consequence of Lemma VIII.B.3 and the representation in Lemma VIII.B.4.

□

Lemma VIII.B.6. *For each $r \in \mathcal{S}^+$, there exists a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $A \in \wp(B(K))$,*

(VIII.9) $$\gamma_r(A) = \left\{ a \in A \mid \sum_{s \in S} a(s)p(s) \geqslant \sum_{s \in S} b(s)p(s) \quad \forall p \in \mathcal{C}_r, \ \forall b \in A \right\}$$

Proof. By VIII-A8 part i. and Lemmas VIII.B.4 and VIII.B.5, $\preceq_r$ is a non-trivial, monotonic, affine, continuous, reflexive, transitive relation, for each $r \in \mathcal{S}^+$. By Ghirardato et al. (2004, Proposition A.2), for every such relation $\preceq_r$, there is a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $a, b \in B$, $a \preceq_r b$ iff

(VIII.10) $$\sum_{s \in S} a(s)p(s) \leqslant \sum_{s \in S} b(s)p(s) \qquad \text{for all } p \in \mathcal{C}_r$$

The required representation follows from Lemma VIII.B.4.

□

Lemma VIII.B.7. *For all $r, s \in \mathcal{S}^+$ with $r \geqslant s$, $\mathcal{C}_s \subseteq \mathcal{C}_r$.*

Proof. By Lemma VIII.C.7 and the representation in Lemma VIII.B.4, for all $r, s \in \mathcal{S}^f$ with $r \geqslant s$, $\preceq_r \subseteq \preceq_s$. The result follows directly by Ghirardato et al. (2004, Proposition A.1). □

Lemma VIII.B.8. *Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on $B(K)$ generated by (VIII.10) with the set of probability measures $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$. Then $\preceq_{\underline{\mathcal{S}}} = \preceq_{\bigcap \mathcal{S}}$.*

Proof. By Lemma VIII.C.7, $\preceq_{\underline{\mathcal{S}}} \supseteq \bigcup_{r\in\mathcal{S}^+} \preceq_r$, and so $\preceq_{\underline{\mathcal{S}}} \supseteq \preceq_{\bigcap \mathcal{S}}$. For the inverse containment, suppose that $a \preceq_{\underline{\mathcal{S}}} b$. If $\{a, b\} \notin \Phi$, the result is immediate, so suppose not. Since $\preceq_{\underline{\mathcal{S}}}$ is affine, by Lemma VIII.B.5, it follows that $x^* \succeq_{\underline{\mathcal{S}}} \alpha(a - b) + x^*$, for any $\alpha \in (0, 1)$, $x \in \Re$ such that $x^*, \alpha(a - b) + x^* \in B(K)$. Take α and x such that this is the case, and let $y = \min_{s\in S}(\alpha(a - b) + x^*)(s)$. We may suppose without loss of generality that $\{x^*, \alpha(a - b) + x^*\} \in \underline{\mathcal{S}}$. Let $f, g, h \in \mathcal{A}$ such that $u \circ f = x^*$, $u \circ g = \alpha(a - b) + x^*$ and $u \circ h = y^*$. By VIII-A9 part ii., since $f \in \gamma(\{f, g\}, \dagger_0)$, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, for every $\beta \in (0, 1)$, there exists $z \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^z(\{f, g\}) = \varnothing$ but $\bar{\gamma}^z(\{f, g_\beta h\}) \neq \varnothing$. It follows from the property i of $\leqq$ and its continuity that, for every $\beta \in (0, 1)$, there exists $r > \underline{\mathcal{S}}$ such that $\gamma_r(\{f, g_\beta h\}) \neq \varnothing$. By Lemma VIII.C.7, the representation in Lemma VIII.B.4 and the monotonicity and transitivity of $\preceq_{\underline{\mathcal{S}}}$, $x^* \succeq_r \beta(\alpha(a - b) + x^*) + (1 - \beta)y^*$, and hence $x^* \succeq_{\bigcap \mathcal{S}} \beta(\alpha(a - b) + x^*) + (1 - \beta)y^*$. Since this holds for every $\beta \in (0, 1)$ and $\preceq_{\bigcap \mathcal{S}}$ is continuous, it follows that $x^* \succeq_{\bigcap \mathcal{S}} \alpha(a - b) + x^*$ and thus, since it is affine, $a \preceq_{\bigcap \mathcal{S}} b$, as required.

□

Lemma VIII.B.9. *The set $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$ is a singleton.*

Proof. By Lemma VIII.B.4, $\preceq_{\underline{\mathcal{S}}}$ is complete. Since, by Lemma VIII.B.8, $\preceq_{\underline{\mathcal{S}}}$ is represented according to (VIII.10) with set of priors $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$, it follows that $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$ is a singleton.

□

Lemmas VIII.C.8 and VIII.C.9 ensure that for all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$ and $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$ whenever r is non-maximal. Lemma VIII.C.10 establishes that, for all $r, s \in \mathcal{S}^+$, if $\mathcal{C}_r \subset \mathcal{C}_s$, then $(\mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r'\in\mathcal{S}^f} \mathcal{C}_{r'}}) = \varnothing$, and similarly for $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$ and $\mathcal{C}_s$ with $s \in \mathcal{S}^+$.

Conclusion of the proof of Lemma VIII.B.2. Define

$$\Xi = \begin{cases} \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r\} & \text{if } \mathcal{S} = \mathcal{S}^f \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \{\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r, \overline{\bigcup_{r\in\mathcal{S}^+} \mathcal{C}_r}\} & \text{if } \mathcal{S} = \mathcal{S}^f \cup \{\overline{\mathcal{S}}\} \end{cases}$$

where the $\mathcal{C}_r$ are as specified in Lemma VIII.B.6, and where, in the second case, $\overline{\mathcal{S}}$ is understood to be a maximal element of $\mathcal{S}$ not belonging to $\mathcal{S}^f$ (ie. a non-full maximal element). It follows from Lemma VIII.B.7 that Ξ is a nested family of sets. Since the $\mathcal{C}_r$ are closed and convex for all $r \in \mathcal{S}^+$ (Lemma VIII.B.6), Ξ is a confidence ranking. By Lemma VIII.B.9, it contains a singleton set, and so is centred. By Lemmas VIII.C.8 and VIII.C.9, Ξ is continuous, and by Lemma VIII.C.10, it is strict.

Define D as follows: for all $A \in \wp(\mathcal{A})$, if $[u \circ A] \in \mathcal{S}^+$, then $D(A) = \mathcal{C}_{[u\circ A]}$, if $A \in \underline{\mathcal{S}}$, then $D(A) = \bigcap_{s\in\mathcal{S}^+} \mathcal{C}_s$, and if $A \in \overline{\mathcal{S}}$, then $D(A) = \overline{\bigcup_{s\in\mathcal{S}^+} \mathcal{C}_s}$. By construction and Lemma VIII.B.7, D respects $\leqq$. Extensionality, Continuity and Φ-Richness of D are immediate from the definition, the fact that D respects $\leqq$, and properties iii, iv and the continuity of $\leqq$. By construction, VIII-A6, and Lemmas VIII.B.6 and VIII.B.8, u, Ξ, D represent the restriction of γ to $\wp(\mathcal{A}) \times \{\dagger_0\}$ according to (VIII.8).

□

Define $c : \Xi \to \Re_{\geqslant 0} \cup \{\infty\}$ as follows. For every $\mathcal{C} \in \Xi$, if $\mathcal{C} = \bigcap_{\mathcal{C}'\in\Xi} \mathcal{C}'$, then $c(\mathcal{C}) = u(\inf \iota(\Upsilon))$; if $\mathcal{C} = \overline{\bigcup_{\mathcal{C}'\in\Xi} \mathcal{C}'}$, then $c(\mathcal{C}) = u(\sup \iota(\Upsilon))$; otherwise, $c(\mathcal{C}) = u(\iota(A))$ for any $A \in D^{-1}(\mathcal{C}) \cap \Upsilon$. (Lemma VIII.B.2, the continuity and Φ-richness of D and the fact that it respects $\leqq$, and properties i and ii of $\leqq$ imply that the value of c in the last case is well-defined: there always exists an appropriate A and the value is independent of the choice of A satisfying the conditions. The fact that u is zeroed implies that c takes non negative values.) By property i of $\leqq$ and the fact that D respects $\leqq$, c is order-preserving and -reflecting with respect to $\subseteq$. Continuity of c follows from the definition and the continuity of $\leqq$, Ξ and D.

Lemma VIII.B.10. *For every $A \in \wp(\mathcal{A})$ and $x > 0$, if $c(D(A)) \leqslant u(x)$, then*

$$\gamma(A, \dagger_x) = \begin{cases} \sup(A, u, D(A)) & \textit{if } \sup(A, u, D(A)) \neq \varnothing \\ \sup(A, u, \{p_\Xi\}) & \textit{if} \sup(A, u, D(A)) = \varnothing \end{cases}$$

where u, Ξ and D are as in Lemmas VIII.C.1 and VIII.B.2.

Proof. Let $A \in \wp(\mathcal{A})$ and $x > 0$ be such that $c(D(A)) \leqslant u(x)$. It follows from the definitions of c and ι, VIII-A4 and VIII-A9 part i. that either $\dagger_0 \notin \gamma(A, \dagger_0)$ or $\dagger_0 \in \gamma(A, \dagger_0)$ and $\dagger_x \notin \gamma(A, \dagger_x)$. It follows from Lemma VIII.B.2 that $\sup(A, u, D(A)) \neq \varnothing$ in the former case and that $\sup(A, u, D(A)) = \varnothing$ in the latter case. Moreover, in the former case, VIII-A4 implies that $\gamma(A, \dagger_{x'}) = \sup(A, u, D(A))$, for all $x' \geqslant 0$, which implies the desired representation. Now consider the latter case. Note that, in this case, $A \in \Phi$.

By the Φ-Richness of D, there exists $\alpha \in (0, 1]$ and $h \in \mathcal{A}$ such that $D(A_\alpha h) = \bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}'$. By the fact that Ξ is centred, $\bigcap_{\mathcal{C}' \in \Xi} \mathcal{C}' = \{p_\Xi\}$. Lemma VIII.B.2 implies that $\gamma(A_\alpha h, \dagger_0) = \sup(A_\alpha h, u, \{p_\Xi\}))$. By VIII-A3, since $\dagger_x \notin \gamma(A, \dagger_x)$, we have that $\gamma(A, \dagger_x)_\alpha h = \gamma(A_\alpha h, \dagger_0) = \sup(A_\alpha h, u, \{p_\Xi\})$. So $\gamma(A, \dagger_x) = \sup(A, u, \{p_\Xi\})$, as required.

□

The desired representation is a direct consequence of Lemmas VIII.B.2 and VIII.B.10, VIII-A6, and the construction of c. Hence the axioms imply the representation, as required.

VIII.B.2 Necessity of axioms

As concerns the (ii) to (i) direction, most cases are straightforward, in the light of the fact that the representation (VIII.2) implies that, if $f \in \bar{\gamma}^x(A, \dagger_0)$, then f is an optimal element in A according to the unanimity rule with the level of confidence $c^{-1}(x)$. So VIII-A1, for example, holds if $f \in \sup(B, u, \mathcal{C})$ implies that $f \in \sup(A, u, \mathcal{C})$, where $\mathcal{C}$ is the confidence level such that $c(\mathcal{C}) = u(x)$; this is evidently the case. Similar considerations apply to VIII-A2 and VIII-A9 part ii. (in the latter case, note that the fact above implies that $mmc(A)$ gives the cost of the highest confidence level which yields an optimal element from A according to the unanimity rule). Apart from the continuity axioms, which shall be discussed below, perhaps the only other axiom requiring comment is VIII-A3. Note that, since the confidence ranking is strict, if $\sum_{s \in S} u(f(s)).p(s) = u(c)$ for all $p \in \mathcal{C}$, for some $\mathcal{C} \in \Xi$ and $f \in \mathcal{A}$, $c \in \Delta(X)$, then for any $\mathcal{C}' \in \Xi$ such that $\sum_{s \in S} u(f(s)).p(s) \neq u(c)$ for some $p \in \mathcal{C}'$, there exists $q, q' \in \mathcal{C}'$ such that $\sum_{s \in S} u(f(s)).q(s) > u(c)$ and $\sum_{s \in S} u(f(s)).q'(s) < u(c)$. Hence, if $\sup(\{f, c\}, u, \mathcal{C}) = \{f, c\}$ for some $\mathcal{C} \in \Xi$, then for every $\mathcal{C}' \in \Xi$, either $\sup(\{f, c\}, u, \mathcal{C}') = \{f, c\}$ or $\sup(\{f, c\}, u, \mathcal{C}') = \varnothing$. It follows, given the properties of the unanimity rule, that for any $A \in \wp(\mathcal{A})$, if $\sup(A, u, \mathcal{C}) = A'$ for some

$\mathcal{C} \in \Xi$, then for every $\mathcal{C}' \in \Xi$, either $\sup(A, u, \mathcal{C}') = A'$ or $\sup(A, u, \mathcal{C}') = \varnothing$. In the light of representation (VIII.2), this implies that VIII-A3 holds.

The final cases of potential interest are VIII-A9, parts i and iii.[34] As concerns the former, consider any $A_n, A \in \wp(\mathcal{A})$, $f_n, f \in \mathcal{A}$ and $x_n, x \in \Re_{\geqslant 0}$, $n \in \mathbb{N}$ with $A_n \to A$, $f_n \to f$ and $x_n \to x$, and such that $f_n \in \gamma(A_n, \dagger_{x_n})$ for all $n \in \mathbb{N}$. We distinguish two cases. If $\dagger_0 \in \gamma(A_n, \dagger_0)$ for all $n \geqslant N$ for some $N \in \mathbb{N}$, then $f \in \gamma(A, \dagger_x)$ follows from the continuity of the EU representation, the continuity of c, the form of representation (VIII.2), and (in the case where $\dagger_0 \notin \gamma(A, \dagger_0)$) the fact, noted above, that $\sup(A, u, \{p_\Xi\}) = \sup(A, u, \mathcal{C})$ for any $\mathcal{C} \in \Xi$ such that $\sup(A, u, \mathcal{C}) \neq \varnothing$. Now consider the case where, for every $N \in \mathbb{N}$, there exists $n > N$ with $\dagger_0 \notin \gamma(A_n, \dagger_0)$. Given the form of representation (VIII.2), it suffices to show that, for any sequences $A_n \to A$ and $f_n \to f$ with $f_n \in \gamma(A_n, \dagger_0)$ for all $n \in \mathbb{N}$, $f \in \gamma(A, \dagger_0)$. Taking such a pair of sequences, we consider the case where there is no $N \in \mathbb{N}$ with $D(A_n) \supseteq D(A)$ for all $n \geqslant N$; the other case is treated similarly. Since there is no such sequence, $D(A) \neq \{p_\Xi\}$, so there exists $B \in \wp(\mathcal{A})$ with $D(B) \subset D(A)$. By the continuity of D, for each such B, there exists $N_B \in \mathbb{N}$ such that $D(A_n) \supseteq D(B)$ for all $n \geqslant N_B$. It follows that, for each $n \geqslant N_B$, $\sum_{s \in S} u(f_n(s)).p(s) \geqslant \sum_{s \in S} u(h(s)).p(s)$ for all $p \in D(B)$ and $h \in A_n$. Hence, by the continuity of the representation, $\sum_{s \in S} u(f(s)).p(s) \geqslant \sum_{s \in S} u(h(s)).p(s)$ for all $p \in D(B)$ and $h \in A$. Since this holds for every B with $D(B) \subset D(A)$, and since the continuity of the confidence ranking and D imply that $D(A) = \overline{\bigcup_{D(B) \subset D(A)} D(B)}$, there cannot be a $q \in D(A)$ such that $\sum_{s \in S} u(f(s)).q(s) < \sum_{s \in S} u(h(s)).q(s)$ for some $h \in A$. So $f \in \sup(A, u, D(A))$; hence $f \in \gamma(A, \dagger_0)$ and so $f \in \gamma(A, \dagger_x)$, as required.

Finally, we consider VIII-A9 part iii. We use the fact, established in Lemma VIII.C.11, that $mmc^{-1}([0, x])$ and $mmc^{-1}([x, \infty))$ are closed in Φ for all $x \in (0, \sup \iota(\Upsilon))$. Suppose that VIII-A9 part iii. does not hold, and that there exists $A \in \Phi$ with $A \in \overline{(\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)} \backslash (\iota^{-1}([0, x]) \cap \Upsilon) \cup (mmc^{-1}([0, x]) \cap \Upsilon^c)$ for some $x \in (0, \sup \iota(\Upsilon))$. It follows that $A \in \overline{(\iota^{-1}([0, x]) \cap \Upsilon)} \backslash (\iota^{-1}([0, x]) \cap \Upsilon)$ or $A \in \overline{(mmc^{-1}([0, x]) \cap \Upsilon^c)} \backslash (mmc^{-1}([0, x]) \cap \Upsilon^c)$. Suppose first that $A \in \overline{(\iota^{-1}([0, x]) \cap \Upsilon)} \backslash (\iota^{-1}([0, x]) \cap \Upsilon)$. If $A \in \Upsilon$, then, by the continuity of D and c and the representation, $A \in \iota^{-1}([0, x]) \cap \Upsilon$. If $A \notin \Upsilon$, then, since $\iota^{-1}([0, x]) \subseteq mmc^{-1}([0, x])$

34. Richness (VIII-A10) is a straightforward consequence of the Φ-Richness of D for part i., and of the representation and the Φ-Richness and continuity of D for part ii. Likewise, VIII-A9 part iv. is a straightforward consequence of the Φ-Richness of D and the form of the representation.

and the latter set is closed by Lemma VIII.C.11, $A \in mmc^{-1}([0,x]) \cap \Upsilon^c$. Now suppose that $A \in \overline{(mmc^{-1}([0,x]) \cap \Upsilon^c)} \backslash (mmc^{-1}([0,x]) \cap \Upsilon^c)$. Since, by the continuity of the representation, Υ^c is closed, $A \in \Upsilon^c$. It follows, since $mmc^{-1}([0,x])$ is closed (by Lemma VIII.C.11), that $A \in mmc^{-1}([0,x]) \cap \Upsilon^c$. All cases contradict the assumption that $A \in \overline{(\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)} \backslash (\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)$; so $(\iota^{-1}([0,x]) \cap \Upsilon) \cup (mmc^{-1}([0,x]) \cap \Upsilon^c)$ is closed in Φ, as required. A similar argument establishes that $(\iota^{-1}([x,\infty]) \cap \Upsilon) \cup (mmc^{-1}([x,\infty]) \cap \Upsilon^c)$ is closed in Φ.

VIII.B.3 Uniqueness

Uniqueness of u follows from the Herstein-Milnor theorem. Without loss of generality, we can thus restrict attention to two representations (u, Ξ, D, c) and (u, Ξ', D', c') of the same choice correspondence γ involving the same utility function. Note firstly that if $\Phi = \varnothing$, then uniqueness of all elements follows from the standard uniqueness properties of the EU representation; henceforth we suppose that $\Phi \neq \varnothing$. (VIII.2) implies that, for any $A \in \Upsilon$, $c(D(A)) = u(\iota(A))$. So the function $c \circ D : \wp(\mathcal{A}) \to \Re_{\geqslant 0} \cup \{\infty\}$ is unique on Υ. By the continuity of c and D, it follows that it is unique on $\overline{\Upsilon}$.

We now consider the uniqueness of D. We first establish that D and D' coincide on Υ. Suppose, for the purposes of reductio, that there exists $\hat{A} \in \Upsilon$ for which they do not coincide. Suppose, without loss of generality, that $p \in D(\hat{A}) \backslash D'(\hat{A})$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in D'(\hat{A})$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ and α can be chosen so that $\alpha \in K$ and $a \in B(K)$. Taking $f \in \mathcal{A}$ such that $u \circ f = a$ and $c \in \Delta(X)$ such that $u(c) = \alpha$, we have that $\sum_{s \in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in D'(\hat{A})$, whereas this is not the case for all $p \in D(\hat{A})$. Let $x \in \Re_{>0}$ be such that $u(x) = c(D(\hat{A})) = c'(D'(\hat{A}))$ (note that since $\hat{A} \in \Upsilon$, it belongs to the domain where $c \circ D$ and $c' \circ D'$ agree). It follows from the representation (VIII.2) and the aforementioned properties of $D'(\hat{A})$ that $f \in \bar{\gamma}^x(\{f, c\})$. However, it follows from the representation and the properties of $D(\hat{A})$ that $f \notin \bar{\gamma}^x(\{f, c\})$, contradicting the assumption that (u, Ξ, D, c) and (u, Ξ', D', c') both represent the choice correspondence γ. Hence D and D' coincide on Υ. Now consider $A \in \overline{\Upsilon} \backslash \Upsilon$, and take any sequence of $A_n \in \Upsilon$ with $A_n \to A$. By the continuity of D, $D(A) = \lim_{n \to \infty} D(A_n)$ and similarly for D'; since D and D' coincide on

Υ, $D(A) = D'(A)$. So D and D' coincide on $\overline{\Upsilon}$.

As concerns the uniqueness of Ξ, we show that $\Xi = \Xi' = D(\overline{\Upsilon})$. It follows from what has just been established that Ξ and Ξ' both contain $D(\overline{\Upsilon})$. We firstly show that there do not exist $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$. Suppose that there does exist such $\mathcal{C} \in \Xi$. Using a separating hyperplane theorem as above, one can construct $f \in \mathcal{A}$, $c \in \Delta(X)$ such that $\sum_{s \in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in \mathcal{C}$, whereas this is not the case for $\bigcap_{A' \in \overline{\Upsilon}} D(A')$. Hence $\{f, c\} \in \Phi$. By the continuity of Ξ, there exists $\mathcal{C}'' \in \Xi$ such that $\mathcal{C}'' \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$ and it is not the case that $\sum_{s \in S} u(f(s))p(s) \geqslant u(c)$ for all $p \in \mathcal{C}''$. By the Φ-Richness and continuity of D, there exists $\alpha \in (0, 1]$, $h \in \mathcal{A}$ such that $D(\{f_\alpha h, c_\alpha h\}) = \mathcal{C}''$. It follows from the representation that $\{f_\alpha h, c_\alpha h\} \in \Upsilon$, and from the order-preserving and -reflecting properties of c that $\iota(\{f_\alpha h, c_\alpha h\}) < \inf \iota(\Upsilon)$, which is a contradiction. So there is no $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \subsetneq \bigcap_{A' \in \overline{\Upsilon}} D(A')$, as required. A similar argument shows that there exists no $\mathcal{C} \in \Xi \backslash D(\overline{\Upsilon})$ with $\mathcal{C} \supsetneq \bigcup_{A' \in \overline{\Upsilon}} D(A')$. Since Ξ is nested, it follows that $\Xi = D(\overline{\Upsilon})$, and similarly for Ξ'. So $\Xi = \Xi'$, as required.

Finally, the uniqueness of c is a direct consequence of the fact that $\Xi = D(\overline{\Upsilon})$, and the uniqueness of D and $c \circ D$ on $\overline{\Upsilon}$.

□

VIII.C Proofs of results mentioned in Appendices VIII.A and VIII.B

Throughout this Appendix, we adopt the notation introduced in Appendix VIII.B.

VIII.C.1 Proofs of results used in Appendix VIII.B

Proof of Theorem VIII.2. Define $\preceq$ by $x \preceq y$ iff $y \in \gamma(\{x, y\})$. Reflexivity is an immediate consequence of $sing$.[35] Transitivity follows from π and α, by the same reasoning as used in the proof of Hill (2009, Theorem 2). It is a straightforward consequence of α that, if

35. Note that there is an error in the statement and proof of Theorem 2 in Hill (2009). The proof (p300) establishes the representation for all non-singleton menus, but not for singleton menus; accordingly, it does not show the reflexivity of the representing relation. This omission is corrected by the addition of the axiom $sing$.

$x \in \gamma(A)$, then $x \geq y$ for all $y \in A$. It remains to show that if $x \geq y$ for all $y \in A$, then $x \in \gamma(A)$.

Suppose that this is not the case: ie. $x \geq y$ for all $y \in A$, but $x \notin \gamma(A)$. We first show that there exists a maximal subset A' of A such that $x \in A'$ and $x \in \gamma(A')$. $sing$ implies that there exists at least one subset such that x is in the image of γ, namely $\{x\}$. The continuity of γ implies that if there were an increasing (under set inclusion) infinite chain A_i of subsets of A such that $x \in \gamma(A_i)$ for all i, then $x \in \gamma(\overline{\bigcup A_i})$ (where $\overline{\bigcup A_i}$ is compact since it is contained in A) so this is a subset containing all the subsets in the chain and such that x is in its image under γ. Hence, by Zorn's Lemma, there exists a maximal subset of A such that $x \in A'$ and $x \in \gamma(A')$. Since $x \notin \gamma(A)$, any such maximal subset cannot be A. Take any such subset A' and consider any $y \in A \backslash A'$. Since $x \in \gamma(A)$ and $x \in \gamma(\{x, y\})$ by hypothesis, π implies that $x \in \gamma(A' \cup \{y\})$, contradicting the maximality of A'. Hence $x \in \gamma(A)$ as required.

The uniqueness is immediate from the definition of $\preceq$. Finally, since π implies Sen's β, the standard theorem for choice correspondences (for example Sen (1971)) implies that, if γ always takes non-empty values then it can be represented by a complete, reflexive, transitive binary relation, without the need for the continuity assumption. □

Lemma VIII.C.1. *There exists a strictly increasing zeroed continuous affine utility function $u : \Delta(X) \to \Re$ representing the restriction of γ to sets of constant acts.*

Proof. VIII-A6 and VIII-A7 imply that $\gamma(A, \dagger_0) \subseteq A$ for all $A \in \wp(\mathcal{A})$ such that $A \subseteq \Delta(X)$. VIII-A1 and VIII-A2 imply that the restriction of $\gamma(\bullet, \dagger_0)$ to sets of constant acts satisfies properties α and π in Theorem VIII.2. By Theorem VIII.2, there exists a complete, transitive, reflexive preference relation $\preceq^0_{|\Delta(X)}$ representing the restriction to $\gamma(\bullet, \dagger_0)$ to sets of constant acts. By VIII-A8 part i., VIII-A3 and VIII-A9 part i., $\preceq^0_{|\Delta(X)}$ is a non-trivial relation satisfying independence and continuity. The existence of an affine u representing $\preceq^0_{|\Delta(X)}$ follows from the Herstein-Milnor theorem. VIII-A9 part i. implies that u is continuous. By VIII-A8 part i., u is strictly increasing. Since u is unique up to positive affine transformation, it can be chosen to be zeroed, as required. By VIII-A4, $\gamma(A, \dagger_x) = \gamma(A, \dagger_0)$ for all $A \in \wp(\mathcal{A})$ with $A \subseteq \Delta(X)$ and all $x \in \Re_{\geqslant 0}$, so u represents the restriction of γ to constant acts, as required.

□

Lemma VIII.C.2. *For all* $A, B \in \wp(\mathcal{A})$ *and* $x \in \Re_{\geqslant 0}$*, if* $A \overset{e.e.}{\simeq} B$ *with relevant correspondence* σ*, then* $\gamma(B, \dagger_x) = \sigma(\gamma(A, \dagger_x))$ *if* $\gamma(A, \dagger_x) \subseteq A$ *and* $\gamma(B, \dagger_x) = \dagger_x$ *if not.*

Proof. Take any A, B and σ with the specified properties. If $\dagger_x \in \gamma(A, \dagger_x)$, then VIII-A8 part iii. and VIII-A6 imply that $\gamma(B, \dagger_x) = \dagger_x$. Now consider the case where $\dagger_x \notin \gamma(A, \dagger_x)$. Let $\bar{A} = \gamma(A, \dagger_x)$; by VIII-A6, $\bar{A} \subseteq \mathcal{A}$. By VIII-A5, there exists $\alpha \in (0, 1]$ and $h \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(A_\alpha h, \dagger_0)$, and by VIII-A3, $\gamma(A_\alpha h, \dagger_0) = \bar{A}_\alpha h$ for any such α and h. For any $f \in \bar{A}$ and $f' \in \sigma(f)$, it follows from VIII-A5 that there exists $\alpha' \in (0, 1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$. For any such α' and h', VIII-A2 (applied with $x = 0$) implies that $\bar{A}_{\alpha'} h' \cup \{f'_{\alpha'} h'\} \subseteq \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, and VIII-A1 implies that $\gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0) = \bar{A}_{\alpha'} h' \cup \{f'_{\alpha'} h'\}$. Similarly, for any $f \in A \backslash \bar{A}$ and $f' \in \sigma(f)$, VIII-A5 implies that there exists $\alpha' \in (0, 1]$ and $h' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$. For any such α' and h', VIII-A2 implies that $\bar{A}_{\alpha'} h' \subseteq \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$ and that if $f'_{\alpha'} h' \in \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, then $f_{\alpha'} h' \in \gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0)$, from which it would follow by VIII-A1 that $f_{\alpha'} h' \in \bar{A}_{\alpha'} h'$, contrary to the assumption. Hence $\gamma((A \cup \{f'\})_{\alpha'} h', \dagger_0) = \bar{A}_{\alpha'} h'$. Repeating this reasoning, we have that, for $\alpha'' \in (0, 1]$ and $h'' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma((A \cup B)_{\alpha''} h'', \dagger_0)$, $\gamma((A \cup B)_{\alpha''} h'', \dagger_0) = \bar{A}_{\alpha''} h'' \cup \sigma(\bar{A})_{\alpha''} h''$. Hence, by VIII-A1, for $\alpha''' \in (0, 1]$ and $h''' \in \mathcal{A}$ such that $\dagger_0 \notin \gamma(B_{\alpha'''} h''', \dagger_0)$, $\gamma(B_{\alpha'''} h''', \dagger_0) = \sigma(\bar{A})_{\alpha'''} h'''$. Since, by VIII-A8 part iii., $\dagger_x \notin \gamma(B, \dagger_x)$, it follows from VIII-A3 that $\gamma(B, \dagger_x) = \sigma(\bar{A})$, as required.

□

Lemma VIII.C.3. *For all* $A \in \Upsilon$*,* $\iota(A) > mmc(A)$*.*

Proof. Note that VIII-A9 part i., VIII-A4 and VIII-A6 imply that, if $\iota(A) = y$, then $\dagger_y \notin \gamma(A, \dagger_y)$. Hence, for all $A \in \wp(\mathcal{A})$, $\iota(\{A_\alpha h \mid \alpha \in (0, 1], h \in \mathcal{A}, A_\alpha h \in \Upsilon\}) \subseteq \{x \in \Re_{\geqslant 0} \mid \bar{\gamma}^x(A) = \varnothing\}$. By VIII-A9 part iv., this set is open, so, for any $A \in \Upsilon$, $\iota(A) > mmc(A)$, as required.

□

Lemma VIII.C.4. *For all* $B \in \wp(\mathcal{A})$*,* $h \in \mathcal{A}$ *and* $\alpha \in (0, 1]$*, if* $B_\alpha h \leqq B$ *and* $\dagger_0 \notin \gamma(B, \dagger_0)$*, then* $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$*. Similarly, if* $B_\alpha h \geqq B$ *and* $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$*, then* $\dagger_0 \notin \gamma(B, \dagger_0)$*.*

Proof. First note that, for any $B \in \mathcal{A}$, $h \in \mathcal{A}$ and $\alpha \in (0, 1]$, $mmc(B_\alpha h) = mmc(B)$ by VIII-A9 part iv. Now suppose that $B \in \mathcal{A}$, $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ are such that $B_\alpha h \leqq B$

and $\dagger_0 \notin \gamma(B, \dagger_0)$. We first show that $\iota(B_\alpha h) = 0$. If this were not the case, then Lemma VIII.C.3 implies that $\iota(B_\alpha h) > mmc(B_\alpha h) = mmc(B)$. It follows, by property ii of $\leqq$ (Lemma VIII.B.1) that $B_\alpha h > B$, contradicting the fact that $B_\alpha h \leqq B$. So $\iota(B_\alpha h) = 0$. By the definition of ι and VIII-A4, there exists $B' \subseteq B$ such that $\gamma(B_\alpha h, \dagger_x) = B'_\alpha h$ for all $x > 0$, whence, by VIII-A9 part i., $\dagger_0 \notin \gamma(B_\alpha h, \dagger_0)$, as required. Similar reasoning establishes the conclusion for the other case.

□

Lemma VIII.C.5. *For every $A, A' \in \wp(B(K))$ such that $A, A' \in \Phi$, there exists $\alpha \in (0, 1]$ and $l \in B(K)$ such that $\alpha A + (1-\alpha)l \equiv A'$.*

Proof. If $A \equiv A'$, then there is nothing to show. Suppose without loss of generality that $A < A'$; the other case is treated similarly. By property iv of $\leqq$ (Lemma VIII.B.1), there exist $\beta \in (0, 1]$ and $l \in \mathcal{A}$ such that $\beta A + (1-\beta)l \geqq A'$. If $\beta A + (1-\beta)l \equiv A'$, then the result has been established; if not, then by continuity of $\leqq$, there exists $\alpha \in (\beta, 1)$ such that $\alpha A + (1-\alpha)l \equiv A'$, as required.

□

Lemma VIII.C.6. *For every $A \in \wp(B(K))$, $l, m \in B(K)$ and $\alpha, \beta \in (0, 1]$ with $\alpha A + (1-\alpha)l, \beta A + (1-\beta)m \in r \in \mathcal{S}^f$, and every $a \in A$, $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$ iff $\beta a + (1-\beta)m \in \gamma(\beta A + (1-\beta)m, \dagger_0)$.*

Proof. The result is an immediate consequence of VIII-A3 if $A \notin \Phi$, so suppose that this is not the case. Without loss of generality, suppose that $\beta \leqslant \alpha$. Consider first the case where $\beta < \alpha$. Note that $\beta A + (1-\beta)m = \frac{\beta}{\alpha}(\alpha A + (1-\alpha)l) + (1 - \frac{\beta}{\alpha})(\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m)$, where $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m \in B(K)$, since it is a $\frac{\alpha\beta-\beta}{\alpha-\beta}$-mix of l and m. Let $B \in \wp(\mathcal{A})$ be such that $\alpha A + (1-\alpha)l = \tilde{u}(B)$ and $h \in \mathcal{A}$ be such that $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m = u \circ h$; so $\beta A + (1-\beta)m = \tilde{u}(B_{\frac{\beta}{\alpha}}h)$. Since $B \equiv B_{\frac{\beta}{\alpha}}h$, Lemma VIII.C.4 implies that $\dagger_0 \notin \gamma(B, \dagger_0)$ iff $\dagger_0 \notin \gamma(B_{\frac{\beta}{\alpha}}h, \dagger_0)$. Hence, by VIII-A3 and VIII-A6, $\gamma(B, \dagger_0)_{\frac{\beta}{\alpha}}h = \gamma(B_{\frac{\beta}{\alpha}}h, \dagger_0)$, which yields the required conclusion.

Now consider the case where $\beta = \alpha$. If $l = m$, the result is immediate, so suppose that $l \neq m$. Note that if there exists $\epsilon \in (0, 1]$ and $k \in B(K)$ with $\epsilon \neq \alpha$ and $\epsilon A + (1-\epsilon)k \equiv \alpha A + (1-\alpha)l$, then, by applying the reasoning in the case above, we get $\alpha a + (1-\alpha)l \in \gamma(\alpha A + (1-\alpha)l, \dagger_0)$ iff $\epsilon a + (1-\epsilon)k \in \gamma(\epsilon A + (1-\epsilon)k, \dagger_0)$ iff $\beta a + (1-\beta)m \in \gamma(\beta A + (1-\beta)m, \dagger_0)$, as required. By the continuity and non-triviality of $\leqq$, $\mathcal{S}^f$ is not a singleton, so there exists $n \in B(K)$ and $\eta \in (0,1)$ such that $\eta(\alpha A + (1-\alpha)l) + (1-\eta)n \not\equiv \alpha A + (1-\alpha)l$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1)$ and $n' \in B(K)$ such that $\delta(\eta(\alpha A + (1-\alpha)l) + (1-\eta)n) + (1-\delta)n' \equiv \alpha A + (1-\alpha)l \in r$. However, $\delta(\eta(\alpha A + (1-\alpha)l) + (1-\eta)n) + (1-\delta)n' = \alpha\eta\delta A + (1-\alpha\eta\delta)(\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l + \frac{1-\eta}{1-\alpha\eta}n) + \frac{1-\delta}{1-\alpha\eta\delta}n')$, with $\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l + \frac{1-\eta}{1-\alpha\eta}n) + \frac{1-\delta}{1-\alpha\eta\delta}n' \in B(K)$ since it is a mix of elements of $B(K)$. So $\alpha\eta\delta$ and $\frac{\delta-\alpha\eta\delta}{1-\alpha\eta\delta}(\frac{\eta-\alpha\eta}{1-\alpha\eta}l + \frac{1-\eta}{1-\alpha\eta}n) + \frac{1-\delta}{1-\alpha\eta\delta}n'$ have the properties required above, and the result is established.

$\square$

Lemma VIII.C.7. *For all $A \in \wp(\mathcal{A})$ and $r, s \in \mathcal{S}^f$ with $r \geqslant s$, if $\gamma_r(A) \neq \varnothing$ then $\gamma_s(A) = \gamma_r(A)$.*

Proof. If $s = r$, there is nothing to show, so suppose not. If $A \notin \Phi$, the result is an immediate consequence of the definition of γ_r, so suppose this is not the case. Without loss of generality, it can be assumed that $A \in r$. (If not, apply the argument below to an $\alpha A + (1-\alpha)l \in r$.) Let $\beta \in (0,1)$ and $m \in B(K)$, be such that $\beta A + (1-\beta)m \in s$ (such β and m exist since $s \in \mathcal{S}^f$). Since, by assumption and VIII-A6, $\dagger_0 \notin \gamma(A, \dagger_0)$, Lemma VIII.C.4 implies that $\dagger_0 \notin \gamma(\beta A + (1-\beta)m, \dagger_0)$; VIII-A3 implies that $\gamma(\beta A + (1-\beta)m, \dagger_0) = \beta\gamma(A, \dagger_0) + (1-\beta)m$, so $\gamma_s(A) = \gamma_r(A)$, as required. $\square$

Lemma VIII.C.8. *For all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$.*

Proof. By Lemma VIII.B.7, $\mathcal{C}_r \supseteq \mathcal{C}_{r'}$ for all $r' < r$. Suppose, for reductio, that $\mathcal{C}_r \supsetneq \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$, so that there exists a point (probability measure) $p \in \mathcal{C}_r \backslash \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) < \alpha \leqslant \phi(q)$ for all $q \in \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0,1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$. Since r is non-minimal in $\mathcal{S}^f$, there exists $l \in B(K)$ and $\beta \in (0,1)$ such that $\beta(\delta\{a, \alpha^*\} + (1-\delta)m) + (1-\beta)l < \delta\{a, \alpha^*\} + (1-\delta)m$. Let $\beta' = \min\{\epsilon \in [\beta, 1] | \ \epsilon(\delta\{a, \alpha^*\} + (1-\delta)m) + (1-\epsilon)l \geqq \delta\{a, \alpha^*\} + (1-\delta)m\}$ (this is a minimum by the continuity of $\leqq$). Taking $f, g, h \in \mathcal{A}$ such that $u \circ f = \delta a + (1-\delta)m$, $u \circ g = \delta\alpha^* + (1-\delta)m$ and $u \circ h = l$, it follows, by the construction, that for any $\eta \in (\beta, \beta')$, $f_\eta h \in \gamma(\{f_\eta h, g_\eta h\}, \dagger_0)$. However, by construction, $f_{\beta'} h \notin \gamma(\{f_{\beta'} h, g_{\beta'} h\}, \dagger_0)$, contradicting VIII-A9 part i. Hence $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. $\square$

Lemma VIII.C.9. *For all non-maximal* $r \in \mathcal{S}^+$, $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.

Proof. By Lemma VIII.B.7, $\mathcal{C}_r \subseteq \mathcal{C}_{r'}$ for all $r' > r$. Suppose, for reductio, that $\mathcal{C}_r \subsetneq \bigcap_{r'>r} \mathcal{C}_{r'}$, so that there exists a point (probability measure) $p \in \bigcap_{r'>r} \mathcal{C}_{r'} \backslash \mathcal{C}_r$. By a separating hyperplane theorem, there is a linear functional ϕ on $ba(S)$, an $\alpha \in \Re$ and an $\epsilon > 0$ such that $\phi(p) \leqslant \alpha - \epsilon$ and $\alpha \leqslant \phi(q)$ for all $q \in \mathcal{C}_r$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s \in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality, α, ϕ and a can be chosen so that $\alpha \in K$, $a \in B(K)$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exists $\delta \in (0, 1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$. Take any $x \in K$ with $x \leqslant \alpha, a(s)$ for all $s \in S$, and let $f, g, h \in \mathcal{A}$ be such that $u \circ f = \delta a + (1-\delta)m$, $u \circ g = \delta\alpha^* + (1-\delta)m$, $u \circ h = \delta x^* + (1-\delta)m$. Let $\beta \in (0, 1)$ be such that $u \circ g_\beta h = \delta(\alpha - \frac{\epsilon}{2})^* + (1-\delta)m$; such a β exists by the choice of a and α and the definition of g and h. By construction, $f \in \gamma(\{f, g\}, \dagger_0)$ and for all $\alpha \in (0, 1]$ and $e \in \mathcal{A}$ such that $\{f, g_\beta h\}_\alpha e > \{f, g\}$, $f_\alpha e \notin \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$. However, $f(s) \in \gamma(\{f(s), h(s)\}, \dagger_0)$ and $g(s) \in \gamma(\{g(s), h(s)\}, \dagger_0)$ for all $s \in S$, and there exists $\alpha \in (0, 1]$ and $e \in \mathcal{A}$ with $\{f, g_\beta h\}_\alpha e > \{f, g\}$. Since, by VIII-A1, VIII-A2, VIII-A3 and VIII-A8 part ii., for all $\alpha \in (0, 1]$ and $e \in \mathcal{A}$, if $\dagger_0 \notin \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$ then $f_\alpha e \in \gamma(\{f, g_\beta h\}_\alpha e, \dagger_0)$, it follows that, for all $\alpha \in (0, 1]$ and $e \in \mathcal{A}$ such that $\{f, g_\beta h\}_\alpha e > \{f, g\}$, $\{f, g_\beta h\}_\alpha e \in \Upsilon$. It follows by properties i, ii and iv of $\leqq$ and its continuity that for all $\sup(\iota(\wp(\mathcal{A}))) > x > mmc(\{f, g\})$, there exists $\alpha \in (0, 1]$, $e \in \mathcal{A}$ with $\iota(\{f, g_\beta h\}_\alpha e) = x$. So for every $x \in \Re_{\geqslant 0}$ such that $\bar{\gamma}^x(\{f, g\}) = \varnothing$, $\bar{\gamma}^x(\{f, g_\beta h\}) = \varnothing$, contradicting VIII-A9 part ii. Hence $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.

□

Lemma VIII.C.10. *For all* $r, s \in \mathcal{S}^+$*, if* $\mathcal{C}_r \subset \mathcal{C}_s$*, then* $(\mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r' \in \mathcal{S}^+} \mathcal{C}_{r'}}) = \varnothing$*. Similarly, for all* $s \in \mathcal{S}^+$, $(\bigcap_{r \in \mathcal{S}^+} \mathcal{C}_r \cap (ri(\mathcal{C}_s))^c) \cap ri(\overline{\bigcup_{r' \in \mathcal{S}^+} \mathcal{C}_{r'}}) = \varnothing$.

Proof. We only consider the case of $r, s \in \mathcal{S}^+$, the other case being treated similarly. Suppose that the condition does not hold, so there exist $r, s \in \mathcal{S}^+$ with $\mathcal{C}_r \subset \mathcal{C}_s$ and $p \in \mathcal{C}_r \cap (ri(\mathcal{C}_s))^c$ but $p \in ri(\overline{\bigcup_{r' \in \mathcal{S}^f} \mathcal{C}_{r'}})$. Let $x_r = \iota(A)$ for any $A \in \Upsilon$ such that $A \in r$ and similarly for x_s; property i of $\leqq$ implies that these are well-defined and Lemma VIII.B.7 implies that $x_s > x_r$. Since S is finite, it follows from a supporting hyperplane theorem (Aliprantis and Border, 2007, Theorem 7.36) that there exists a linear functional ϕ supporting $\mathcal{C}_s$ at p; ie. such that $\phi(q) \geqslant \phi(p)$ for all $q \in \mathcal{C}_s$. Let $\phi(p) = \alpha$. Since B is finite-dimensional, there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s\in S} a(s)q(s)$ for any $q \in ba(S)$. Without loss of generality ϕ can be chosen so that $a \in B(K)$ and $\alpha \in K$. By construction, $\{a, \alpha^*\} \in \Phi$. Since $r \in \mathcal{S}^f$, there exist $\delta \in (0, 1]$ and $m \in B(K)$ such that $\delta\{a, \alpha^*\} + (1-\delta)m \in r$. Take $f, g \in \mathcal{A}$ such that $u \circ f = \delta a + (1-\delta)m$ and $u(g) = \delta\alpha^* + (1-\delta)m$. Since, by construction, $\delta a + (1-\delta)m \in \gamma_s(\delta\{a, \alpha^*\} + (1-\delta)m)$, it follows that $mmc(\{f, g\}) \geqslant x_s > x_r$. Finally, it follows from the construction that for any $z > \alpha$ and $g' \in \mathcal{A}$ such that $u(g') = \delta z^* + (1-\delta)m$, and any $\beta \in (0, 1)$, $\dagger_0 \in \gamma(\{f_\alpha h, (g_\beta g')_\alpha h\}, \dagger_0)$ for all $\alpha \in (0, 1], h \in \mathcal{A}$ such that $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f, g\}$. Since, whenever $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f, g\}$, $\dagger_{x_r} \notin \gamma(\{f_\alpha h, (g_\beta g')_\alpha h\}, \dagger_{x_r})$, it follows that each $\{f_\alpha h, (g_\beta g')_\alpha h\}$ with $\{f_\alpha h, (g_\beta g')_\alpha h\} \equiv \{f, g\}$ belongs to $\iota^{-1}([0, x_r]) \cap \Upsilon$. By the continuity of $\leqq$, every neighborhood of $\{f, g\}$ contains such a $\{f_\alpha h, (g_\beta g')_\alpha h\}$; so $\{f, g\}$ is an element of Φ on the boundary of $(\iota^{-1}([0, x_r]) \cap \Upsilon) \cup (mmc^{-1}([0, x_r] \cap \Upsilon^c)$, but not belonging to this set, contradicting VIII-A9 part iii. So there exist no such r, s, as required.

□

Lemma VIII.C.11. *Let γ be represented according to (VIII.2). Then, for all $x \in (0, \sup\iota(\Upsilon))$, $mmc^{-1}([0, x])$ and $mmc^{-1}([x, \infty))$ are closed in Φ.*

Proof. For ease of presentation, we adopt the following notation: for any $x \in \Re$, $\mathcal{C}_x = c^{-1}(x)$. Take $x \in (0, \sup\iota(\Upsilon))$; to show that $mmc^{-1}([0, x])$ is closed, suppose not, and suppose that $A \in \left(\overline{mmc^{-1}([0, x])}\backslash mmc^{-1}([0, x])\right) \cap \Phi$. Let $A_n \in mmc^{-1}([0, x])$ be a sequence with $A_n \to A$. It follows from representation (VIII.2) and the definition of mmc that, for every $n \in \mathbb{N}$, every $f_n \in \sup(A_n, u, \{p_\Xi\})$ and every $x' > x$, there exists $p, q \in \mathcal{C}_{x'}$ and $g, h \in A_n$ such that $\sum_{s\in S} u(f_n(s)).p(s) < \sum_{s\in S} u(h(s)).p(s)$ and $\sum_{s\in S} u(f_n(s)).q(s) > \sum_{s\in S} u(g(s)).q(s)$. It follows from the representation and the fact that $A \in \left(\overline{mmc^{-1}([0, x])}\backslash mmc^{-1}([0, x])\right) \cap \Phi$ that for some $x' > x$ and all $f \in \sup(A, u, \{p_\Xi\})$, $\sum_{s\in S} u(f(s)).p(s) \geqslant \sum_{s\in S} u(h(s)).p(s)$ for all $p \in \mathcal{C}_{x'}$ and all $h \in A$. Since $A_n \to A$, it follows from the continuity of the representation that the hyperplane $u \circ f - u \circ h = 0$ in $\Delta(S)$ supports $\mathcal{C}_{x'}$, for at least one $h \in A$ and $f \in \sup(A, u, \{p_\Xi\})$. If this holds for some $x' > x$, it must also hold for any $x' > x'' > x$ by the same reasoning; it follows from the fact that Ξ is a nested family that there exists a support point of $\mathcal{C}_{x'}$ that is also a support point of $\mathcal{C}_{x''}$. Moreover, since $A \in \Phi$, such a point is not a support point of $\overline{\bigcup_{\mathcal{C}'\in\Xi} \mathcal{C}'}$. Since support points are those that do not belong to the relative interiors of the appropriate sets, it follows that $(C_{x''} \cap (ri(C_{x'}))^c) \cap ri(\overline{\bigcup_{\mathcal{C}'\in\Xi} \mathcal{C}'}) \neq \varnothing$, contradicting the strictness of Ξ. So $mmc^{-1}([0,x])$ is closed in Φ. The closeness of $mmc^{-1}([x,\infty])$ in Φ is a straightforward consequence of the continuity of the unanimity representation.

□

VIII.C.2 Proofs of results in Appendix VIII.A

Proof of Proposition VIII.A.1. The 'if' direction is straightforward. The 'only if' direction is a simple corollary of the proof of Theorem VIII.1. On the one hand, if γ^1 and γ^2 are confidence equivalent, they yield identical choices over menus consisting entirely of constant acts; hence the utilities are the same up to positive affine transformation. On the other hand, if they are confidence equivalent, the family of functions $\{\gamma_r | r \in \mathcal{S}^f\}$ defined in the proof of Theorem VIII.1 are the same, and so the confidence rankings are the same. □

Proof of Proposition VIII.A.2. Let the assumptions of the Proposition be satisfied and let (u, Ξ, D_1, c_1) and (u, Ξ, D_2, c_2) represent γ^1 and γ^2 respectively. Beyond using some terminology introduced in Section VIII.3.2.3, the following notation shall prove useful. For any $A \in \wp(\mathcal{A})$, define $\mathcal{C}_A \in \Xi$ as follows: $\mathcal{C}_A = \max\{\mathcal{C} \in \Xi \mid \sup(A, u, \mathcal{C}) \neq \varnothing\}$. We consider the three parts in turn.

Part (i). The right-to-left implication is straightforward, so we only consider the left-to-right direction. It follows from representation (VIII.2) that $c_1(D_1(A)) = u(\iota_1(A))$ for all $A \in \Upsilon_1$ and similarly for decision maker 2. Since γ^1 is more decision averse than γ^2, $\iota_2(A) \leqslant \iota_1(A)$ for all $A \in \wp(\mathcal{A})$, and hence in particular $A \in \Upsilon_2$ implies $A \in \Upsilon_1$. It follows that, for each $A \in \Upsilon_2$, $c_1(D_1(A)) \geqslant c_2(D_2(A))$. This property extends to $\overline{\Upsilon_2}$ by the continuity of the D and c.

Define $D_3 : \wp(\mathcal{A}) \to \Xi$ by $D_3(A) = D_2(A)$ whenever $A \in \overline{\Upsilon_2}$ and $D_3(A) = \min\{D_1(A), D_2(A)\}$ otherwise. By definition, and the fact that for $A \in \overline{\Upsilon_2}$, $c_1(D_1(A)) \geqslant c_2(D_2(A))$, $c_2 \circ D_3$ satisfies the ordering conditions with respect to $c_1 \circ D_1$ in the Proposition. We now show that it is a cautiousness coefficient and that (u, Ξ, D_3, c_2) represents γ^2. Continuity of D_3 off the boundary of Υ_2 follows from the continuity of D_1, D_2 and the minimum. To establish continuity for menus on the boundary, it suffices to show that for every $A \in \overline{\Upsilon_2} \backslash \Upsilon_2$, $D_2(A) \subseteq D_1(A)$. For every $A \in \overline{\Upsilon_2} \backslash \Upsilon_2$, it follows from representation (VIII.2) that $D_2(A) = \mathcal{C}_A$; since, as noted above, $\Upsilon_2 \subseteq \Upsilon_1$, then either $A \in \Upsilon_1$, in which case $D_1(A) \supset D_2(A)$ by the definition of Υ, or $A \in \overline{\Upsilon_1} \backslash \Upsilon_1$, in which case $D_1(A) = \mathcal{C}_A = D_2(A)$. So $D_2(A) \subseteq D_1(A)$, and D_3 satisfies continuity. Extensionality and Φ-richness of D_3 follow from the extensionality and Φ-richness of D_1 and D_2, so D_3 is a cautiousness coefficient. Finally, we show that (u, Ξ, D_3, c_2) represents γ_2. Since, by definition, $D_3(A) = D_2(A)$ for all $A \in \overline{\Upsilon_2}$, $c_2 \circ D_3 = c_2 \circ D_2$ on $\overline{\Upsilon_2}$. Moreover, for all $A \notin \overline{\Upsilon_2}$, $\sup(A, u, D_3(A)) \neq \varnothing$, and so, by representation (VIII.2) (and the strictness of the confidence ranking), $\sup(A, u, D_3(A)) = \sup(A, u, D_2(A))$. Since (u, Ξ, D_2, c_2) represents γ_2, so does (u, Ξ, D_3, c_2), as required.

Part (ii). By the representation (VIII.2), for any $A \in \Phi_1$, $mmc_1(A) = c_1(\mathcal{C}_A)$ and similarly for decision maker 2. Moreover that, by a simple supporting hyperplane argument, for every $\mathcal{C} \in \Xi \backslash \{\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}\}$, there exists $A' \in \Phi_1$ such that $\mathcal{C}_{A'} = \mathcal{C}$. (It suffices to consider any supporting hyperplane of $\mathcal{C}$, and take the menu containing an act and a constant act corresponding to the hyperplane, as in Lemmas VIII.C.8–VIII.C.10, for example.) Note finally that, since $A \in \Phi_1$ iff $\sup(A, u, \overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}) \neq \varnothing$, it follows from the fact that γ^1 and γ^2 are represented by the same u and Ξ that $\Phi_1 = \Phi_2$. It follows from representation (VIII.2) that γ^1 is less cost motivated than γ^2 iff $mmc_1(A) \geqslant mmc_2(A)$ for all $A \in \Phi_1$. By the first remark above, this holds iff $c_1(\mathcal{C}_A) \geqslant c_2(\mathcal{C}_A)$ for all $A \in \Phi_1$; by the second remark, this holds iff $c_1(\mathcal{C}) \geqslant c_2(\mathcal{C})$ for all $\mathcal{C} \in \Xi \backslash \{\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}\}$. By the continuity of c_1 and c_2, this holds iff $c_1(\mathcal{C}) \geqslant c_2(\mathcal{C})$ for all $\mathcal{C} \in \Xi$, yielding the required equivalence.

Part (iii). By the representation (VIII.2), for any $A \in \Upsilon_1$, $\iota_1(A) = c_1(D_1(A))$ and similarly for decision maker 2. By representation (VIII.2), γ^1 is more motivation-calibrated decision averse than γ^2 iff for all $A, B \in \wp(\mathcal{A})$, $\iota_1(A) \leqslant mmc_1(B) \Rightarrow \iota_2(A) \leqslant mmc_2(B)$. Given the observations made in the proof of part (ii), this holds iff, for all $A \in \Upsilon_2$ and $\mathcal{C} \in \Xi \backslash \{\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}\}$, $c_1(D_1(A)) \leqslant c_1(\mathcal{C}) \Rightarrow c_2(D_2(A)) \leqslant c_2(\mathcal{C})$. Since c_1 and c_2 are order-preserving and -reflecting, this holds iff $D_1(A) \subseteq \mathcal{C} \Rightarrow D_2(A) \subseteq \mathcal{C}$ for all $A \in \Upsilon_2$ and $\mathcal{C} \in \Xi \backslash \{\overline{\bigcup_{\mathcal{C}' \in \Xi} \mathcal{C}'}\}$, and hence iff $D_1(A) \supseteq D_2(A)$ for all $A \in \Upsilon_2$. So γ^1 is more motivation-calibrated decision averse than γ^2 iff $D_1(A) \supseteq D_2(A)$ for all $A \in \Upsilon_2$. Note moreover that if γ^1 is more motivation-calibrated decision averse than γ^2, then $\Upsilon_2 \subseteq \Upsilon_1$: for any $A \in \Upsilon_2$, $D_2(A) \supset \mathcal{C}_A$, and hence, since γ^1 is more motivation-calibrated decision averse, $D_1(A) \supset \mathcal{C}_A$, so $A \in \Upsilon_1$.

It follows immediately from representation (VIII.2) and the uniqueness conditions in Theorem VIII.1 that the existence of a cautiousness coefficient satisfying the conditions in part (iii) implies that γ^1 is more motivation-calibrated decision averse than γ^2. To show the left-to-right direction of part (iii), we show that if $\Upsilon_2 \subseteq \Upsilon_1$ and $D_1(A) \supseteq D_2(A)$

for all $A \in \Upsilon_2$, then there exists cautiousness coefficient D_3 representing γ^2 such that $D_3(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$. To this end, let $D_3(A) = \min\{D_1(A), D_2(A)\}$ for all $\mathcal{A}$. This is a cautiousness coefficient: extensionality, continuity and Φ-richness follow from the extensionality, continuity and Φ-richness of D_1 and D_2, the continuity of the minimum, and the fact that $D_1(A) \supseteq D_2(A)$ on Υ_2. By definition, $D_3(A) \subseteq D_1(A)$ for all $A \in \wp(\mathcal{A})$. By a similar argument to that used in the proof of part (i), (u, Ξ, D_3, c_2) represents γ^2. So (u, Ξ, D_1, c_1) and (u, Ξ, D_3, c_2) satisfy the required conditions.

□

Proposition VIII.C.3. *Let γ^1 and γ^2 satisfy axioms VIII-A1–VIII-A10 and be confidence equivalent.*

— *If does not follow from γ^1 being more decision averse and less cost motivated than γ^2 that γ^1 is more motivation-calibrated decision averse than γ^2.*

— *It does not follow from γ^1 being more decision averse and more motivation-calibrated decision averse than γ^2 that γ^1 is less cost motivated than γ^2.*

Proof. Consider any (u, Ξ, D_2, c_2) satisfying the conditions in Theorem VIII.1 and let γ^2 be the choice correspondence it represents. Take two continuous functions $f_D : c_2(\Xi) \to c_2(\Xi)$ and $f_c : c_2(\Xi) \to \Re$ such that f_D is surjective and f_c is strictly increasing. Define $D_1 : \wp(\mathcal{A}) \to \Xi$ by $D_1(A) = c_2^{-1}(f_D(c_2(D_2(A))))$ and $c_1 : \Xi \to \Re_{\geqslant 0}$ by $c_1(\mathcal{C}) = f_c(c_2(\mathcal{C}))$. It is straightforward to check that these are well-defined cautiousness coefficients and cost functions respectively; let γ^1 be the choice correspondence represented by (u, Ξ, D_1, c_1). Let us say that a function $\Gamma : X \to Y$, where $X \subseteq Y$, is *upper-valued* if, for every $x \in X$, $f(x) \geqslant x$. Note that, by Proposition VIII.A.2 (and the order-reflecting and preserving properties of the cost function), γ^1 and γ^2 are ordered by decision aversion iff $f_c \circ f_D$ is upper-valued, they are ordered by motivation-calibrated decision aversion iff f_D is upper-valued, and they are ordered by cost motivation iff f_c is upper-valued.

To establish the first part of the proposition, it thus suffices to find f_c and f_D such that f_c is upper-valued, f_D is not, but $f_c \circ f_D$ is. Let $a = \inf c_2(\Xi)$ and $b = \sup c_2(\Xi)$, and consider $f_D(x) = \frac{1}{b-a}(x-a)^2 + a$, and $f_c(x) = x + \frac{b-a}{2}$. It is straightforward to check that f_D is not upper-valued, whilst f_c and $f_c \circ f_D$ are, yielding the required example. Similarly, taking for example $f_D(x) = (b-a)^{\frac{1}{2}}(x-a)^{\frac{1}{2}} + a$, and $f_c(x) = \frac{3}{4}(x-a) + a$ for $x \leqslant \frac{b-a}{4} + a$, and $f_c(x) = \frac{13}{12}(x-a) - \frac{b-a}{12} + a$ for $x \geqslant \frac{b-a}{4} + a$ yields an example establishing the second part of the proposition.

□

Bibliography

Akerlof, G. A. (1991). Procrastination and Obedience. *The American Economic Review*, 81(2):1–19.

Aliprantis, C. D. and Border, K. C. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition.

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Bewley, T. F. (1986 / 2002). Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25(2):79–110.

Billingsley, P. (2009). *Convergence of Probability Measures*. John Wiley & Sons.

Buturak, G. and Evren, O. (2014). A Theory of Choice When "No Choice" is an Option. Technical report, NES.

Danan, E. (2003a). A behavioral model of individual welfare. Technical report, Université Paris 1.

Danan, E. (2003b). Revealed Cognitive Preference Theory. Technical report, EUREQua, Université de Paris 1.

Dekel, E., Lipman, B. L., and Rustichini, A. (2001). Representing Preferences with a Unique Subjective State Space. *Econometrica*, 69(4):891–934.

Dhar, R. (1997). Consumer Preference for a No-Choice Option. *The Journal of Consumer Research*, 24(2):215–231.

Ergin, H. and Sarver, T. (2010). A unique costly contemplation representation. *Econometrica*, 78(4):1285–1339.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2):133–173.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Gruber, M. J. (1996). Another puzzle: The growth in actively managed mutual funds. *The Journal of Finance*, 51(3):783–810.

Gul, F. and Pesendorfer, W. (2001). Temptation and Self-Control. *Econometrica*, 69(6):1403–1435.

Hart, K. P., Nagata, J.-i., and Vaughan, J. E. (2004). *Encyclopedia of general topology*. Elsevier.

Hill, B. (2009). Confidence in preferences. *Social Choice and Welfare*, 39(2):273–302.

Hill, B. (2013). Confidence and decision. *Games and Economic Behavior*, 82:675–692.

Hill, B. (2014). Incomplete preferences and Confidence. Technical report, HEC Paris.

Holmstrom, B. (1984). On the Theory of Delegation. In Boyer, M. and Kihlstrom, R., editors, *Bayesian Models in Economic Theory*. North-Holland, New York.

Kelley, J. L. (1975). *General Topology*. Springer.

Koopmans, C. (1964). On the flexibility of future preferences. In Shelley, M. and Bryan, G., editors, *Human judgements and rationality*. John Wiley and Sons.

Kopylov, I. (2009). Choice deferral and ambiguity aversion. *Theoretical Economics*, 4(2):199–225.

Kreps, D. M. (1979). A Representation Theorem for 'Preference for Flexibility'. *Econometrica*, 47:565–576.

Malkiel, B. G. (1995). Returns from investing in equity mutual funds 1971 to 1991. *The Journal of Finance*, 50(2):549–572.

Marschak, J. and Miyasawa, K. (1968). Economic Comparability of Information Systems. *International Economic Review*, 9(2):137–174.

Masatlioglu, Y. and Ok, E. A. (2005). Rational choice with status quo bias. *Journal of Economic Theory*, 121(1):1–29.

O'Donoghue, T. and Rabin, M. (2001). Choice and Procrastination. *Quart. J. Econ.*, 116(1):121–160.

Ortoleva, P. (2010). Status quo bias, multiple priors and uncertainty aversion. *Games and Economic Behavior*, 69(2):411–424.

Ortoleva, P. (2013). The price of flexibility: Towards a theory of Thinking Aversion. *Journal of Economic Theory*, 148(3):903–934.

Riella, G. and Teper, R. (2014). Probabilistic dominance and status quo bias. *Games and Economic Behavior*, 87:288–304.

Savochkin, A. (2014). Mistake Aversion and a Theory of Robust Decision Making. Technical report, Collegio Carlo Alberto.

Sen, A. K. (1971). Choice Functions and Revealed Preference. *Rev. Econ. Stud.*, 38(3):307–317.

Stigler, G. J. (1961). The Economics of Information. *The Journal of Political Economy*, 69(3):213–225.

Tversky, A. and Shafir, E. (1992). Choice Under Conflict: The Dynamics of Deferred Decision. *Psychological Science (Wiley-Blackwell)*, 3(6):358–361.

IX Dynamic Consistency and Ambiguity: A Reappraisal

Abstract

Dynamic consistency demands that a decision maker's preferences over contingent plans agree with his preferences in the planned-for contingency. What counts are the contingencies the decision maker envisages – and plans for – rather than contingencies selected by a theorist. We show how this simple point resolves some crucial purported difficulties for ambiguity models in dynamic settings. Firstly, when properly formulated, dynamic consistency is compatible with consequentialism and non-expected utility. Secondly, the perspective provides a principled justification for the restrictions to certain families of beliefs (à la Epstein and Schneider (2003)) in applications to dynamic choice problems. Thirdly, the value of information under non-expected utility is revealed to be non-negative as long as the information offered does not compromise information that the decision maker had otherwise expected to receive. We also give behavioral foundations for the contingencies the decision maker envisages, in the form of a representation theorem. [1]

Keywords: Decision under Uncertainty; Dynamic Consistency; Dynamic Choice; Value of Information; Epistemic Contingency.

JEL classification: D81, D83, D90.

1. This is a slightly modified version of HEC Paris Research Paper No ECO/SCD-2013-983.

IX.1 Introduction

IX.1.1 Motivation

One of the principal challenges in, and to, the literature on non-expected utility models for decision under uncertainty is that posed by application in dynamic situations. Violation of expected utility purportedly brings with it some problematic consequences, which have been at the base of criticism of the whole literature on ambiguity (for some examples, see Al Najjar et al. (2009) and discussions in Epstein and Le Breton (1993), Ghirardato (2002)). Putting aside subtleties in the terminology and definitions, the gist of the problem lies in the argument that consequentialism – the decision maker ignores sunk costs (or the history in the decision tree) when deciding at any node – and dynamic consistency – the decision maker's preferences over contingent plans agree with his preferences in the planned-for contingency – are incompatible with non-expected utility. Given the *primae facie* attraction of these properties of dynamic choice, not to mention their importance for economic applications, the option of abandoning one of them to leave space for non-expected utility is unappetizing. All the worse, some have concluded, for non-expected utility theories.

The main thesis of this paper is that this generally accepted state-of-play rests on a mistake. Standard formalizations of the dynamic consistency principle use sets of contingencies that are effectively imposed by the theorist. However, since the principle involves the decision maker's plans, it is the contingencies that he himself envisages – and plans for – that are relevant under any reasonable version of the principle. Once the principle is formulated with the contingencies the decision maker himself envisages, the apparent incompatibility with non-expected utility is resolved.

This apparently innocuous point has wide-ranging conceptual and economic ramifications. First of all, it implies a reconceptualization in terms of what the decision maker envisages about his future states of belief. Dynamic consistency, as considered in the literature on decision under uncertainty, is revealed to be essentially a condition on the relationship between two sorts of belief: the decision maker's current beliefs about the state of the world and his beliefs about what he will believe in the future about the state of the world.

Secondly, this relationship turns out to be particularly strong in certain cases: so strong, in fact, that one can draw conclusions, just on the basis of the decision maker's current beliefs, about what he thinks he might believe in the future, and hence about the contingencies he envisages. In particular, the standard modelling assumption that the decision maker is using the same decision tree as the theorist implies that his ex ante beliefs must be of a certain form. In perhaps the most significant economic contribution of the paper, this observation provides a principled justification for the restriction to a specific family of beliefs in the application of non-expected utility models to dynamic choice problems. In such problems, the use of a set of beliefs not in this family implies that the decision maker do not think that he is facing the decision tree used by the theorist.

Thirdly, the approach provides a new perspective on the purported propensity of non-expected decision makers to turn down free information. Information about the decision tree that the decision maker is faced with may impact upon his beliefs about his future beliefs, and hence, in the light of the implications mentioned above, his current beliefs. Incorporating this factor into his choice of whether to accept free information, it turns out that he will always accept an offer of free information, as long as it does not compromise information he had otherwise expected to receive.

In the rest of the Introduction, we present the notion of dynamic consistency, explain our challenge to the standard argument against non-expected utility, and set out the main contributions and plan of the paper, as well as some relations to existing literature.

IX.1.2 Dynamic consistency, events and epistemic contingencies

To present the notion of dynamic consistency for decision under uncertainty, let us adopt the standard setup in the literature, where a decision maker has ex post and ex ante preferences over acts – functions from a given ('objective') set of states of the world, called the state space and understood to represent all payoff-relevant factors, to a set of consequences. We focus on the dynamic consistency of preferences, rather than the dynamic consistency of behavior (see for example Strotz (1955); Karni and Safra (1989, 1990); Siniscalchi (2009, 2011)). Put succinctly, dynamic consistency demands harmony between ex ante preferences over contingent plans and preferences after the realisation of the planned-for contingency: the decision maker's ex post preferences correspond to his ex ante preferences over plans involving the contingency in question. At first glance, application of this principle would seem to require determining which contingencies the decision maker envisages holding, and hence plans for. Somewhat surprisingly, there is little discussion of this question in the literature. Rather, it is generally assumed that one can associate to any pair of acts two contingencies – one in which the acts necessarily yield the same consequence, and one in which this is not the case. The idea is that the acts can be thought of as plans for the contingency in which they differ. In this way, one avoids the question of which contingencies the decision maker himself envisages; the contingencies are, so to speak, imposed by the theorist in the pairs of acts considered.

An important consequence of this approach is that the contingencies are events, that is, sets of states of the world. It is clear on reflection that, strictly speaking, the contingencies that the decision maker plans for cannot simply be events, because they must factor in the fact that the decision maker knows them to hold. A decision maker does not plan for the contingency that event A holds but he does not know it, and it makes no sense to speak of his preference or state of belief after having learnt that such a contingency has arose, for that is not something that he can learn. Taking contingencies as events is shorthand: the contingency is not strictly speaking the fact that the event in question has occurred, but rather the fact that the decision maker learns that the event has occurred. Hence planned-for contingencies are more accurately represented not by events, but rather by future possible states of knowledge or belief (or by the preferences generated by these states of belief). We call these *epistemic contingencies*.

This pedantry would be of no interest if it did not undermine the practice described above of considering contingencies to be events on which acts yield the same or different consequences. Although, as we have seen, any event has a natural epistemic contingency corresponding to it (namely, the contingency that the decision maker learns only that the event in question has occurred), the converse is not true. A contingency in which the decision maker alters his probability for a certain event without giving it probability 0 or 1 (for example, by some form of non-Bayesian learning) is an epistemic contingency which cannot be naturally associated to an event. More importantly, what counts as overlapping, disjointness and exhaustivity for contingencies differs between events and epistemic contingencies. Whilst a (non-empty) event A and the event containing the whole state space S are not exclusive, the epistemic contingency in which the decision maker learns only that A and that in which he learns nothing (ie. where he learns S) are: they cannot both concurrently occur. This has important consequences when one considers which contingencies may need to be planned for. While, when reasoning with events, planning for the

Figure IX.1 – Dynamic consistency in the standard (static) Ellsberg urn (values in dollars)

	R	B	Y
f_1	10	0	10
g_1	0	10	10
f_2	10	0	0
g_2	0	10	0

eventualities A and A^c exhausts all possibilities, this is not the case when the contingencies are epistemic: learning only that A, learning only that A^c and learning nothing are three distinct epistemic contingencies, each of which occurs without any of the others occurring. Hence, even considering the decision maker's preference between acts that only differ on the event A, it does not follow that the only contingencies that he envisages – and plans for – are necessarily (learning that) A and (learning that) A^c: he could equally well envisage, and plan for, the contingency in which he learns nothing (or, indeed, the contingencies in which he learns only B or only B^c for a different event B).

The set of epistemic contingencies that the decision maker envisages at a particular moment forms a decision tree; we shall call this the decision maker's *subjective* tree. The set of events that the theorist considers also forms a tree; since these trees can be defined independently of considerations regarding the decision maker's opinion on the relevant possible future states, we shall call these *objective* trees. As noted above, the class of objective trees can be embedded into the class of subjective trees, in the sense that for each objective tree, there could exist a decision maker who envisages all and only the (epistemic versions of) the events in that tree. However, there are subjective trees that correspond to no objective tree. For such trees, the argument that consequentialism and dynamic consistency rule out non-expected utility behavior does not go through. [2]

To illustrate this point, let us consider a standard example of violation of expected utility for decision under uncertainty, the Ellsberg one-urn example. Consider a decision maker who is told that an urn contains ninety balls, thirty of which are red (R), and the

2. It might be objected that if the decision maker's envisaged future states of belief – the 'subjective information structure' he is using, if you will – are relevant, then they should be represented explicitly in the state space. As discussed in Section IX.2.2 (in particular, Remark IX.2), a setup involving such an extended state space is equivalent to the one used here, and all the points made below continue to hold when reformulated in these terms. See also Appendix IX.A.

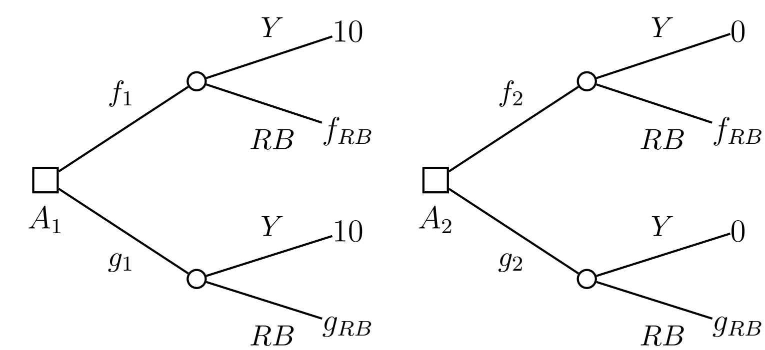

(a) Ex ante choice situations

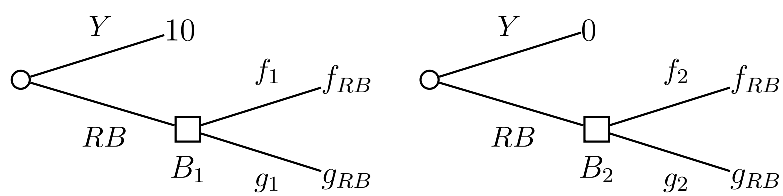

(b) Ex post choice situations

Figure IX.2 – Dynamic consistency in the standard (dynamic) Ellsberg urn

rest of which are black (B) and yellow (Y) in an unknown proportion, and is asked for his preferences over the bets on the colour of a ball drawn from the urn shown in Figure IX.1. The standard preference pattern is $f_1 \prec g_1$ and $f_2 \succ g_2$, violating the sure-thing principle (Savage's P2 (1954)).

A standard formulation of the argument against these Ellsberg preferences using the principles of consequentialism and dynamic consistency is as follows. One considers not only the decision maker's preferences over the bets in Figure IX.1, but one also imagines that, after the ball is drawn, the decision maker receives information as to whether the ball is yellow or not (and no other information), and is asked again for his preferences over the bets. Using f_{RB} to refer to the partial act, defined on the event $\{R, B\}$,[3] and coinciding with f_1 and f_2 on these states, and similarly for g_{RB}, one can, the story goes, represent the decision maker's choices before and after the information is revealed by the trees in Figure IX.2. Consequentialism demands that all that counts are the consequences of one's choices; hence, in particular, it demands that the choices made at the B nodes in Figure IX.2(b) are independent of the values that could have been won if the ball were yellow (ie. the upper branches of the trees). So the choices made at B_1 and B_2 are the same. Dynamic

3. A partial act is a partial function from states to consequences, defined on a subset of the state space.

consistency in preferences asks for coherence between one's preferences over contingent plans concerning future options and one's preferences in the appropriate contingencies. So, if the decision maker prefers f_1 to g_1 at node A_1 in Figure IX.2(a), then, since the only difference between the consequences yielded by these bets occurs when the ball is not yellow, this can be thought of as a preference over plans for the contingency that the ball is not yellow. Dynamic consistency thus demands that he prefer f_1 to g_1 at node B_1 in Figure IX.2(b) (after the resolution of the uncertainty about whether the ball is yellow or not). Similar reasoning applies to the right hand trees (nodes A_2 and B_2) and the acts f_2 and g_2. It follows that, contrary to the Ellsberg preferences, the decision maker must have the same preferences over f_2 and g_2 as over f_1 and g_1.

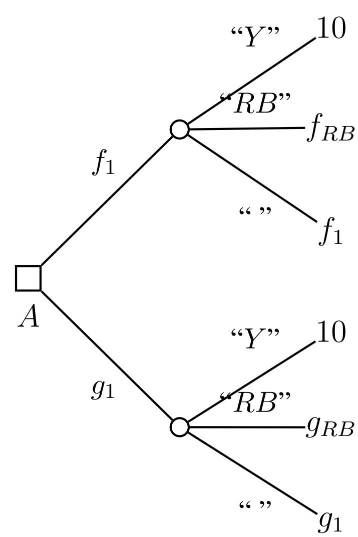

(a) Ex ante choice situation

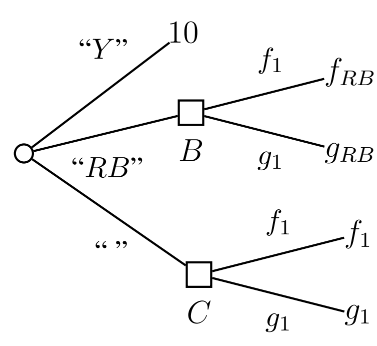

(b) Ex post choice situation

Figure IX.3 – Dynamic consistency in the non-standard (dynamic) Ellsberg urn (choice between f_1 and g_1)

This argument is most naturally read as involving objective trees, especially given that the trees in Figure IX.2 are motivated by the choice of acts considered, independently of considerations about the contingencies the decision maker himself envisages. Certainly, if these trees were read as involving epistemic contingencies, then they involve a non-trivial assumption: that the decision maker envisages only two future states of belief – one in which he learns (only) that the ball is yellow and one in which he learns (only) that it is not. This assumption is far from innocent: it does not hold, for example, for a decision maker who envisages that he might learn whether the ball is yellow or not, but also that he might learn nothing.[4] In this case, the appropriate (subjective) decision trees are not those in Figure IX.2 but rather those given in Figure IX.3 (we only give the trees for the choice between f_1 and g_1; the trees for the choice between f_2 and g_2 are similar). There are three branches corresponding to the epistemic contingency in which the decision maker learns that the ball is yellow (denoted here with inverted commas to emphasise the difference from mere events), the epistemic contingency in which he learns that the ball is not yellow, and the contingency in which he learns nothing at all (denoted " ": the message is empty). On these trees, dynamic consistency can no longer be applied as in the argument presented above. It applied non-trivially on the trees in Figure IX.2 because the choice at the initial node in the ex ante case (node A_1) corresponds to the choice at a single node (B_1) in the ex post case, and hence can be interpreted as a plan for that node. By contrast, in Figure IX.3, the choice between the acts f_1 and g_1 at node A does not correspond to the choice at a single node in the ex post case; thus it cannot be thought of as a plan for a single ex post contingency. Hence dynamic consistency does not necessarily apply – for example, if the preferences over f_1 and g_1 differ at nodes B and C in Figure IX.3(b), then dynamic consistency does not imply anything about the choice in Figure IX.3(a). So the argument against Ellsberg-style behavior no longer goes through. This illustrates a general point that will be made precise in the sequel: consequentialism, non-expected utility preferences and dynamic consistency are not necessarily incompatible, if dynamic consistency is defined using subjective trees, with the contingencies the decision maker does in fact envisage.

4. To give another example, it also does not hold for a decision maker who thinks that he might learn either whether the ball is yellow or not, or whether he has won his bet or not.

IX.1.3 Outline of the paper

This paper has three aims. The first is to develop the rebuttal of the standard argument from dynamic consistency and consequentialism to expected utility in the context of decision under uncertainty. The second, and economically most important aim is to confront the proposed perspective with some of the toughest purported problems for non-expected utility models in dynamic contexts, such as application in dynamic choice problems and the question of value of information. The final aim is to fully defend the central notions of the proposal, and in particular that of epistemic contingency.

A first step will be to propose a formal representation of epistemic contingencies. Assuming, as is standard in this literature, that the decision maker's utility function remains fixed, we shall take contingencies to be states of belief, interpreted as the future states of belief that the decision considers possible ex ante. Although most of the main points continue to hold for other non-expected utility models, for concreteness we shall work with the maxmin expected utility model (Gilboa and Schmeidler, 1989),[5] and accordingly take epistemic contingencies to be sets of priors. We thus implicitly assume that the decision maker forms preferences according to the same decision rule in all contingencies as well as ex ante (this assumption has been explicitly discussed in Eichberger and Kelsey (1996); Sarin and Wakker (1998), for example).

We then go on to formulate a notion of dynamic consistency in terms of these epistemic contingencies, and show that it is possible in general for non-expected utility decision makers to satisfy this condition in conjunction with consequentialism; the point illustrated in the example above does indeed hold in general. The main conceptual contribution of the paper lies in the observation that dynamic consistency in the context of decision under uncertainty demands harmony between two sorts of belief: the beliefs the decision maker currently has about the state of the world (which determine his current preferences, including his preferences over plans), and the beliefs he has about what he will believe about the states of the world (which correspond to the possible future preferences he envisages). As an illustration, we characterise the constraints the decision maker's anticipations about his future beliefs place on his current beliefs, under dynamic consistency.

We then turn to the consequences of this refined version of dynamic consistency for

5. It seems that versions of all results except the representation theorem (Theorem IX.1) are relatively easy to obtain for many other non-expected utility models.

economic applications. In standard applications to dynamic situations, it is often implicitly assumed that the decision maker knows the decision tree with which he is faced, and that this tree is as described by the theorist. In the current setup, these are cases where the decision maker's subjective tree corresponds to an objective one. We show that this particular form of the subjective decision tree has consequences for the decision maker's current beliefs, if he is to satisfy dynamic consistency with respect to epistemic contingencies. It implies that his ex ante set of priors satisfies a 'factorisability' property that is the equivalent in our framework of Epstein and Schneider's (2003) rectangularity. It is sometimes suggested that the need for a restriction to such rectangular sets of priors is a limitation of the Epstein and Schneider (2003) approach. Our result can be interpreted as showing that this criticism is unfounded: the specific shape of the set of priors is a consequence of the standard assumption that the decision maker is using the same decision tree as the theorist. This non-arbitrary vindication of the use of a restricted family of sets of priors of the sort suggested by Epstein and Schneider (2003) is perhaps the most significant economic contribution of the paper. Since this is the key to treating sequential choice problems, it shows that the approach developed here provides a solid, well-defended account of dynamic choice under non-expected utility.

Another pressing question concerning dynamic choice concerns the value of information, which is standardly argued to be possibly negative for non-expected utility decision makers (Al Najjar et al., 2009; Siniscalchi, 2011). The proposed approach implies a reconceptualisation of the issue, insofar as the decision maker's choice to accept an offer of information that he had not previously anticipated obtaining entails a change in his envisaged future states of belief. Given the relationship between envisaged future beliefs and current beliefs mentioned above, it follows that the latter may also change on learning that he will obtain the information; indeed, this change will occur even before the information is actually received. To the extent that this change is a result of his choice to accept the information, the situation has the flavour of a moral hazard problem. Incorporating these elements, we show that, when properly defined, the value of information is non-negative whenever the information offered does not compromise the reception of other information that the decision maker was previously expecting to receive.

An important difference between the objective and subjective trees on which the preceding considerations are based is the question of foundations. Whereas objective trees are part of the setup, as posed by the theorists, subjective trees incorporate the contingencies that the decision maker envisages, and hence encapsulate elements of his state of mind. However, without some connection to observable choice behavior (beyond the apparent violation of dynamic consistency), the notion of epistemic contingency, and the previous considerations building on it, are on shaky ground. The final contribution of the paper is to treat this issue, by providing behavioral foundations for this notion.

The paper is organised as follows. Section IX.2 describes the framework and the representation of epistemic contingencies used. In Section IX.3, the notion of dynamic consistency on subjective trees is defined, and immediate consequences, including the possibility of satisfying dynamic consistency in tandem with consequentialism and non-expected utility, are explored. Section IX.4 examines the case of a decision maker whose subjective tree corresponds to an objective one, and the consequences for the application of non-expected utility models to dynamic problems. Section IX.5 analyses the consequences of the perspective developed here for the question of value of information. The final section, Section IX.6, provides choice-theoretical foundations for the notion of epistemic contingency that underlies the proposed approach. Proofs are collected in an Appendix.

IX.1.4 Related Literature

The literature on dynamic consistency is too large to be cited in its entirety. A thorough discussion of the question and the literature on dynamic arguments for expected utility under risk is given in Machina (1989); papers showing or discussing the inconsistency between dynamic consistency, consequentialism and non-expected utility in the case of uncertainty include Hammond (1988); Epstein and Le Breton (1993); Ghirardato (2002). Of the papers introducing update rules for non-expected utility models or considering dynamic choice in these models, the closest is without doubt Epstein and Schneider (2003). Indeed, as explained in Section IX.4, the notion of $\mathcal{P}$-rectangularity introduced here is essentially a version of their rectangularity condition adapted to our framework, and there are evident technical similarities between Proposition IX.4 and their results. However, as explained in the subsequent discussion, the conceptual contributions are different, so much so that our result can be read as a vindication of their restriction on sets of priors. Their approach has been adopted with other non-expected utility models (for example Maccheroni et al. (2006); Klibanoff et al. (2009)), and the perspective developed here applies similarly.

Conceptually, the closest suggestion to the one proposed here that we have been able to find was made in Gilboa et al. (2009), where it was suggested that some of the events required to exhibit violations of Savage's Sure Thing Principle (P2) in some of the Ellsberg examples are 'highly contrived' and 'will never be observed by the decision maker'. However, as they remark, this point does not hold for the Ellsberg one-urn example that we considered in Section IX.1.2. Indeed, the approach proposed here focusses on what the decision maker expects to learn, rather than what he can learn or observe.

The argument that non-expected utility decision makers may exhibit aversion for information has been proposed by Wakker (1988) for the case of decision under risk (see also Machina (1989) and references within); see Eichberger et al. (2007); Al Najjar et al. (2009) for an analogous argument for the case of decision under uncertainty. Siniscalchi (2011) contains a thorough and insightful discussion of the issue, though arrives at a diagnosis that is clearly very different from ours.[6] For Siniscalchi, the nub of the matter lies in the trade-off between information acquisition and commitment: if commitment is more valuable than information acquisition, then the decision maker may turn down the information. The approach set out here, by contrast, draws on the fact that the choice to accept an offer of information will affect one's envisaged possible future states of belief. If this change involves a loss of some other information that the decision maker had previously expected to receive, then he may turn down the offer.

As discussed in Section IX.6, the representation result proposed there can be understood (technically) as a contribution to the literature on preference for flexibility or decision making under unforeseen contingencies initiated by Kreps (1979, 1992); Dekel et al. (2001). The closest result in that literature to the current one is doubtless Epstein et al. (2007, Theorem 2), insofar as both our representation (IX.5) and their representation (8) involve a set – of utilities in their case, of priors in ours – in each contingency, and use a maxmin rule. However, they use an expected utility 'aggregator' to form ex ante preferences, where we use a dominance aggregator (see Remark IX.7), and they use a richer framework (and, accordingly, different proof methods), involving lotteries over menus. Other related approaches (though not in the menu framework) include Mukerji (1997); Ghirardato (2001).

Other relations (notably technical ones) with existing literature are mentioned at appropriate points in the discussion.

6. Machina (1989) attributes a similar approach to the issue for decision under risk to Edi Karni.

IX.2 Preliminaries

IX.2.1 Setup and technical notions

We use a version of the standard Anscombe-Aumann framework (Anscombe and Aumann, 1963), as adapted by Fishburn (1970). Let S be a non-empty finite set of *states*, and let Σ be the set of subsets of S, called *events*. $\Delta(\Sigma)$ is the set of probability measures on (S, Σ). Where necessary, we use the Euclidean topology on $\Delta(\Sigma)$. X is a nonempty set of *outcomes* endowed with a metric under which it is compact; a *consequence* is a Borel probability measure on X. $\Delta(X)$ is the set of consequences; where necessary, we use the weak convergence topology, under which this space is compact metric (Billingsley, 2009, p 72). *Acts* are functions from states to consequences; $\mathcal{A}$ is the set of acts, endowed with the product topology, under which it is compact metric. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in \Re$, $0 \leqslant \alpha \leqslant 1$, the mixture $\alpha f + (1-\alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s,x) = \alpha f(s,x) + (1-\alpha)h(s,x)$. The mixture relation is extended to sets of acts and acts pointwise: for $A \subseteq \mathcal{A}$, $h \in \mathcal{A}$ and $0 \leqslant \alpha \leqslant 1$, $\alpha A + (1-\alpha)h = \{\alpha f + (1-\alpha)h \mid f \in A\}$. We write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$ and similarly for $A_\alpha h$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$. We use $\wp$ to denote the set of closed non-empty subsets of; hence, in particular, $\wp(\mathcal{A})$ is the set of closed non-empty subsets of $\mathcal{A}$. Where required, we use the Hausdorff topology on $\wp(\mathcal{A})$ (see for example Aliprantis and Border (2007, Section 3.17)). For any $A \in \wp(\mathcal{A})$, $conv(A)$ is the set of finite mixtures of elements of A: $conv(A) = \{\sum_{i=1}^{n} \alpha_i f_i | \ \alpha_i \in [0,1]$ with $\sum_{i=1}^{n} \alpha_i = 1,\ f_i \in A\}$. It is obvious that $conv(A) \in \wp(\mathcal{A})$.

The symbol $\preceq$ (potentially with subscripts) will be used to denote preference relations over $\mathcal{A}$; as standard $\prec$ and $\sim$ denote the asymmetric and symmetric parts of $\preceq$. We adopt the standard notion of null event with respect to a preference relation $\preceq$: an event $A \subseteq S$ is $\preceq$-null iff $f \sim g$ whenever $f(s) = g(s)$ for all $s \in A^c$. Throughout, we shall use $\preceq$ to denote the decision maker's current preference relation.

Finally, we introduce some standard properties of functions $J : \Re^n \to \Re$. J is: *constant additive* if $J(\mathbf{x} + a\mathbf{e}) = J(\mathbf{x}) + a$ for all $\mathbf{x} \in \Re^n$, $a \in \Re$ and where $\mathbf{e}$ is the unit vector; *positively homogeneous* if $J(\alpha\mathbf{x}) = \alpha J(\mathbf{x})$ for all $\mathbf{x} \in \Re^n$, $\alpha \geqslant 0$; *monotonic* if $J(\mathbf{x}) \geqslant$

$J(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \Re^n$ such that $x_i \geqslant y_i$ for all $i \in \{1, \ldots, n\}$, and *strongly monotonic* if it is monotonic and $J(\mathbf{x}) > J(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \Re^n$ such that $x_i \geqslant y_i$ for all $i \in \{1, \ldots, n\}$ with strict inequality for some i.

IX.2.2 Epistemic contingencies

Until Section IX.5, we shall work under the assumption that the set of epistemic contingencies the decision maker envisages is given; in Section IX.6 we do away with this assumption and give a choice-theoretic foundation for this set. Moreover, as is standard, we shall implicitly assume a situation where there are two time periods: the current (ex ante) one and the future (ex post) one. In accordance with a significant number of non-expected utility models (Gilboa and Schmeidler, 1989; Ghirardato et al., 2004), we represent the decision maker's beliefs about the state of the world by sets of priors, that is convex, closed subsets of $\Delta(\Sigma)$.[7] The closed convex set $\mathcal{C} \subseteq \Delta(\Sigma)$ represents the decision maker's current beliefs. The set of closed convex sets $\{\mathcal{C}_i\}_{i \in I}$, where I is a finite index set, represents the future beliefs the decision maker envisages himself as possibly having. These can be thought of as the epistemic contingencies the decision maker anticipates possibly being in. To focus on the issue of dynamic (in)consistency related to non-expected utility, we assume, as standard, that the decision maker uses the same affine utility function $u : \Delta(X) \rightarrow \Re$ ex ante and ex post. Finally, we assume that the decision maker's preferences (both ex ante and ex post) can be represented according to the maxmin expected utility rule (Gilboa and Schmeidler, 1989). According to this rule, preferences $\preceq'$ corresponding to a set of priors $\mathcal{C}'$ are such that, for every $f, g \in \mathcal{A}$, $f \preceq' g$ iff

$$\min_{p \in \mathcal{C}'} \sum_{s \in S} u(f(s))p(s) \leqslant \min_{p \in \mathcal{C}'} \sum_{s \in S} u(g(s))p(s) \qquad \text{(IX.1)}$$

Note that expected utility preferences correspond to the special case where the set of priors is a singleton.

It follows from these assumptions that $\preceq$, the decision maker's current preference relation, satisfies (IX.1) with respect to $\mathcal{C}$. Let $\preceq_i$ be the preference relation on $\mathcal{A}$ satisfying

7. For simplicity, we abstract from the debate about the degree to which, given the lack of separability of ambiguity and ambiguity attitude in some of these models, the sets of priors can be thought of as (purely) beliefs.

(IX.1) with respect to $\mathcal{C}_i$ (for all $i \in I$). Note that, under the aforementioned assumptions, using the set of envisaged ex post beliefs $\{\mathcal{C}_i\}_{i \in I}$ to represent the set of epistemic contingencies is equivalent to using a representation by the set of ex post preferences $\{\preceq_i\}_{i \in I}$, with each preference relation satisfying the Gilboa and Schmeidler (1989) axioms. For ease of presentation, we shall often use the representation in terms of preferences in the next few sections. Note that, whilst $\preceq$ can be taken as an observable primitive, the set $\{\preceq_i\}_{i \in I}$ cannot.

Remark IX.1. Consequentialism states that the decision maker's ex post or conditional preference does not depend on 'other branches' in the decision tree. The standard formulation, which says that preferences conditional on an event A do not depend on what the acts deliver on A^c (Ghirardato, 2002), does not capture consequentialism in the current framework, since the $\preceq_i$ do not necessarily correspond to updating on events. Nevertheless consequentialism is implied by the fact that the ex post preferences are determined entirely by the set of priors and the utility function: if other parts of the tree had an influence on the preferences under this contingency, then (IX.1) would not be an appropriate representation of these preferences. Given that it is tacitly assumed in the framework, we shall not explicitly mention the condition in subsequent discussion.

Remark IX.2. An alternative setup would involve expanding the state space so that the decision maker's ex post beliefs are explicitly represented in the states. As discussed in Appendix IX.A, under the appropriate assumptions, this framework is essentially equivalent to the one proposed above; moreover, all the points made below could equivalently have been in this 'extended state space' framework. Similar remarks apply to other possible setups, such as those involving a subjective state space in the style of Kreps (1992); Dekel et al. (2001) or a space of 'signals' or an information structure, beyond the objective state space S. See also Remark IX.6 (in Section IX.5) on the relation to information structures, and Remark IX.7 (in Section IX.6) on the relation to the literature on subjective state spaces.

IX.3 Dynamic consistency on subjective trees

In Section IX.1.2, we saw on a simple example how dynamic consistency on subjective trees, involving epistemic contingencies, does not imply expected utility preferences, even in the presence of consequentialism. In the current section, this point is made more rigorously: we formulate a version of dynamic consistency for subjective trees, and give an example to show that dynamic consistency, consequentialism and non-expected utility ex ante preferences are compatible, for all sets of anticipated ex post preferences (that is, for all subjective decision trees). In order to explore what dynamic consistency implies for non-expected utility models, we shall also briefly consider its consequences for the relationship between ex ante and ex post preferences in the case of maxmin expected utility preferences.

IX.3.1 Introducing Dynamic Consistency

To introduce the definition of dynamic consistency over subjective trees, it is reasonable to begin with the standard dynamic consistency condition over objective trees considered in the literature. Despite considerable differences between authors, the following condition seems to be fairly representative.

Standard Dynamic Consistency (SDC). For every $f, g \in \mathcal{A}$ and partition $\{A_j\}_{j \in J}$ of S, if $f \preceq_{A_j} g$ for every $\preceq$-non-null A_j, then $f \preceq g$, and moreover, if any of the $\preceq_{A_j}$ orderings are strict, then so is the $\preceq$ one.

Here $\preceq_{A_i}$ are the preferences conditional on A_i. Another standard condition used in the literature (see for example Ghirardato (2002)) is the special case of this condition applied to two-element partitions and acts which coincide on one element of the partition. It is straightforward to show that, given transitivity of preferences, this condition is equivalent to the one given above.

Standard Dynamic Consistency (SDC) captures the idea that when faced with any objective decision tree, comprising of a partition of events, if the decision maker will always prefer one act to another in all of the future eventualities, then this is the case under his current preferences. This corresponds to the requirement that his ex ante preferences should be coherent with his ex post preferences in the relevant contingencies, under the assumption that the relevant contingencies are the events of a partition. As argued previously, this assumption does not hold in general, and when it does not, then the ex post preferences that count are those in the (epistemic) contingencies that the decision maker does in fact envisage ex ante – that is the set of anticipated future preferences $\preceq_i$ – rather than those in the contingencies corresponding to the events imposed by the theorist – the $\preceq_{A_j}$. It is straightforward to extend SDC to incorporate epistemic contingencies, and subjective trees.

Dynamic Consistency (DC). For all $f, g \in \mathcal{A}$, if $f \preceq_i g$ for all $\preceq_i$, then $f \preceq g$, and moreover, if any of the $\preceq_i$ orderings are strict, then so is the $\preceq$ one.

This is a straightforward modification of Standard Dynamic Consistency, where the preferences conditional on events in a partition are replaced by the anticipated future preferences. Note that the non-nullness condition is not needed, because the set of preferences $\{\preceq_i \mid i \in I\}$ is assumed to contain only the future preferences that the decision maker envisages, and already leaves out future contingencies to which he gives 'probability zero'. This is the refined notion of dynamic consistency that we propose, and with which we shall work throughout the rest of this paper.

Remark IX.3. Whilst SDC involves universal quantification over objective decision trees (in the form of quantification over the partitions of the state space), DC does not involve universal quantification over subjective decision trees (ie. over all possible sets of epistemic contingencies). One reason for this is that there is no normative justification for such quantification, least of all from the (English-language formulation of the) dynamic consistency principle itself. The principle is concerned with the decision maker's preferences over contingent plans, and these involve the contingencies he actually envisages, not a set of contingencies that he *could have* but in fact does not envisage. Indeed, since he does not envisage such counterfactual sets of contingencies, there is no need for him to plan for them, so dynamic consistency has nothing to say about them. Moreover, adding quantification over sets of epistemic contingencies to DC would yield an empty condition, which can never be satisfied.[8] Finally, the apparent analogy between a quantified version of DC and SDC may be an artefact of the setup adopted here; as shown in Appendix IX.A, when SDC is formulated appropriately in an extended state space with epistemic contingencies represented in the states, it turns out to be equivalent to DC.

8. A simple way to see this is to consider two sets $\{\preceq_i\}_{i\in I}$, $\{\preceq_j\}_{j\in J}$ of epistemic contingencies, where an event A is null with respect to each preference relation in $\{\preceq_i\}_{i\in I}$ and A^c is null with respect to each preference relation in $\{\preceq_j\}_{j\in J}$. No ex ante preference relation can satisfy DC with respect to both of these sets.

IX.3.2 Dynamic Consistency and Non-expected utility

As explained in the Introduction, a classic argument shows that Standard Dynamic Consistency and Consequentialism (in tandem with other basic assumptions on preferences) implies that the ex ante preferences are represented according to expected utility. This no longer goes through when SDC is replaced by Dynamic Consistency, which makes reference to epistemic contingencies rather than partitions of events. Indeed, no matter what epistemic contingencies the decision maker envisages, there exists a non-expected utility ex ante preference that satisfies Dynamic Consistency with respect to this set of contingencies (recall from Remark IX.1 that consequentialism is automatically satisfied in the framework used here).

Proposition IX.1. *For any set $\mathcal{D}$ of probability measures on I with $p(i) > 0$ for all $i \in I$ and $p \in \mathcal{D}$, let $\preceq_{\mathcal{D}}$ be the preference relation represented according to (IX.1) with the set of priors $\mathcal{C}_{\mathcal{D}} = \{\sum_{i \in I} p(i).q_i |\, p \in \mathcal{D},\ q_i \in C_i\}$. Then $\preceq_{\mathcal{D}}$ and $\{\preceq_i\}_{i \in I}$ satisfy Dynamic Consistency.*

This can be thought of as a possibility result, showing that once one recognises and adopts the contingencies which the decision maker himself envisages, dynamic consistency ceases to be inconsistent with non-expected utility theory. Rather, for any non-singleton set of anticipated ex post preferences (ie. for any non-trivial subjective decision tree), there is a non-expected utility ex ante preference satisfying dynamic consistency with respect to them: it suffices to take $\mathcal{C}_{\mathcal{D}}$ for any non-singleton $\mathcal{D}$.

IX.3.3 Respecting Dynamic Consistency

If Dynamic Consistency does not imply expected utility, then one might wonder what conditions it does place on ex ante preferences. Conceptually, this is basically a question concerning the relationship between two sorts of beliefs the decision maker has: his current beliefs about the state of the world – which determine his ex ante preferences – and his beliefs about what he will believe in the future about the state of the world – which correspond to the epistemic contingencies he envisages. What constraints do what the decision maker thinks about his future possible states of belief place on his current beliefs?

Technically, the question is analogous to the problem of aggregating non-expected utility decision makers' preferences. Indeed, dynamic consistency is evidently a strong dominance or Pareto condition on potential future preferences with respect to current preferences, of the sort commonly used in the literature on aggregation. Moreover, as noted by Crès et al. (2011), forming ex ante preferences with the set of priors given in Proposition IX.1 is equivalent to using a maxmin expected utility functional as the 'aggregator' of ex post maxmin EU preferences into ex ante ones.[9] However, as has been noted in literature on aggregating sets of priors, dynamic consistency does not imply that the decision maker must aggregate his future preferences according to this maxmin EU aggregator. (Crès et al. (2011, Section 2.3) give an example where this is not the case.) In fact, the aggregator may be more general, albeit not very illuminating, as the following proposition shows.

Proposition IX.2. *$\{\preceq_i\}_{i\in I}$ and $\preceq$ satisfy Dynamic Consistency if and only if there exists a constant additive, positively homogeneous, monotonic function $\phi : \Re^I \to \Re$ that is strongly monotonic on $Ra((\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(\cdot)p(s))_{i\in I})$ such that $\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s) = \phi((\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(f(s))p(s))_{i\in I})$ for all $f \in \mathcal{A}$.*

Here $Ra((\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(\cdot)p(s))_{i\in I}) \subseteq \Re^I$ is the range of the $(\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(\cdot)p(s))_{i\in I}$. Ghirardato et al. (2004) provide a representation of constant additive, positively homogeneous and monotonic functionals in terms of a generalised α-maxmin expected utility representation, and this could be substituted into this proposition without changing its meaning.

Whilst technically useful below, this result is conceptually not particularly enlightening. It does have the virtue of showing that the refined version of Dynamic Consistency proposed here is not empty: it does impose non-trivial constraints on one's current beliefs ($\mathcal{C}$) given one's beliefs about one's future beliefs ($\{\mathcal{C}_i\}_{i\in I}$). Whilst, for largely technical reasons, there seems to be little more that can be said in general about the relationship between the ex ante set of priors and the ex post sets of priors under DC, in certain specific cases the relationship may be stronger and more interesting, as we shall see in the next section.

Remark IX.4. One can find supplementary conditions beyond DC that characterize the maxmin EU aggregator used in Proposition IX.1. One such condition draws its intuition from the observation that although there are cases where maxmin EU decision makers can satisfy dynamic consistency without taking a maxmin EU aggregator of their anticipated

9. Their Propostion 1 shows that using the set of priors $\mathcal{C}_{\mathcal{D}} = \{\sum_{i\in I}p(i)q_i |\ p\in\mathcal{D},\ q_i\in C_i\}$ to form maxmin EU preferences over acts is equivalent to evaluating an act by its maxmin expected ex post valuation, taken with the set of priors $\mathcal{D}$, over the set of ex post functionals.

future preferences, such ways of aggregating future preferences are fragile in the sense that, if the set of ex post preferences happened to be different, the aggregated current preferences will generally not belong to the maxmin EU family. Hence the following notion of *robustness* of an aggregator.

Definition IX.1. An aggregator $\phi : \Re^I \to \Re$ is *robust* if, for every set of maxmin EU preferences $\{\preceq_i\}_{i \in I}$, the derived preferences $\preceq$, defined by $f \preceq g$ iff $\phi((\min_{p \in \mathcal{C}_i} \sum_{s \in S} u(f(s))p(s))_{i \in I}) \leqslant \phi((\min_{p \in \mathcal{C}_i} \sum_{s \in S} u(g(s))p(s))_{i \in I})$ for all $f, g \in \mathcal{A}$, is a maxmin EU preference.

A rule for aggregating one's anticipated future preferences into a current preference is robust if, irrespective of what one's future preferences are, whenever they are maxmin EU one is guaranteed to have maxmin EU current preferences. This notion is all that needs to be added to DC to characterize the maxmin EU aggregator.

Proposition IX.3. *ϕ is a robust aggregator satisfying DC iff there exists a unique set $\mathcal{D}$ of probability measures on I with $p(i) > 0$ for all $i \in I$ and such that* $\phi(\mathbf{x}) = \min_{p \in \mathcal{D}} \sum_{i \in I} p(i) x_i$.

IX.4 Dynamic consistency and dynamic choice problems

We now turn to some economic consequences of the refined notion of dynamic consistency. Economists are often interested in cases where the decision maker is told which decision tree he is facing, and where this is the (objective) decision tree used by the theorist. At first glance, these cases would seem to pose a challenge for the approach developed above, given its reliance on the notion of subjective decision tree. Consider, for example, a maxmin EU decision maker considering the Ellsberg urn, with set of priors $\{p|\ p(B) \in [0, \frac{2}{3}],\ p(R) = \frac{1}{3}\}$ (where B is the event that the black ball is drawn and R is the event that the red ball is drawn). Suppose that he is told that he will be faced with the 'Dynamic Ellsberg' decision tree given in Section IX.1.2 (Figure IX.2(a)). This decision maker cannot satisfy both dynamic consistency and consequentialism: herein lies the problem, and, it seems, the introduction of a refined notion of dynamic consistency involving epistemic contingencies does nothing to alleviate this difficulty. Moreover, in many dynamic problems, such as planning problems, intertemporal portfolio problems and applications to saving, these are precisely the sorts of implications of non-expected utility models that inhibit the use of standard techniques, such as backwards induction (or 'folding back') methods. In this section, we examine to what extent the proposed approach can handle such cases.

From the perspective of the framework introduced above, the situations just described concern decision makers whose subjective decision trees corresponds to objective ones. An objective decision tree involves as contingencies the events in a (finite) partition $\mathcal{P}$ of the state space. In the situations described, each of the decision maker's epistemic contingencies $\{\preceq_i\}_{i\in I}$ corresponds to learning exactly one of these events (and no two epistemic contingencies correspond to the same event). Hence, the subjective tree corresponds to the objective tree with partition $\mathcal{P}$ precisely when, for each $i \in I$, there exists exactly one element A_i of $\mathcal{P}$ such that A_i^c is $\preceq_i$-null, and for any $i, j \in I$, if $i \neq j$, then $A_i \neq A_j$. If this holds for a set of epistemic contingencies $\{\preceq_i\}_{i\in I}$ and a partition $\mathcal{P}$, we say that $\{\preceq_i\}_{i\in I}$ is $\mathcal{P}$*-ontic*.

As the following result shows, under this assumption, dynamic consistency places a strong constraint on the form of the decision maker's ex ante set of priors.

Proposition IX.4. *Let $\{\preceq_i\}_{i\in I}$ and $\preceq$ satisfy DC. Suppose moreover that, for some partition $\mathcal{P}$, $\{\preceq_i\}_{i\in I}$ are $\mathcal{P}$-ontic. Then there exists a unique set $\mathcal{D}$ of probability functions on I with $p(i) > 0$ for all $i \in I$ and $p \in \mathcal{D}$, such that $\mathcal{C} = \{\sum_{i\in I} p(i).q_i | \, p \in \mathcal{D}, q_i \in \mathcal{C}_i\}$.*

For a partition $\mathcal{P}$ and a set of priors $\mathcal{C}$, we say that $\mathcal{C}$ is $\mathcal{P}$*-rectangular* if for every $A_j \in \mathcal{P}$ there exists a set $\mathcal{C}_j$ of probability measures with support in A_j, and there exists a set $\mathcal{C}_0$ of probability measures on $\mathcal{P}$ such that $\mathcal{C} = \{\sum_{A_j\in P} p(A_j).q \mid p \in \mathcal{C}_0, \; q \in \mathcal{C}_j\}$. A $\mathcal{P}$-rectangular set of priors has a particular 'shape': it can be 'factorized' into a set of priors over $\mathcal{P}$ multiplied by sets of priors on each element of the partition. This may be thought of as the equivalent in the present setup of the notion of rectangularity defined by Epstein and Schneider (2003) .

Proposition IX.4 thus tells us that, under the assumption of Dynamic Consistency, whenever the decision maker's subjective decision tree corresponds to an objective tree (it is $\mathcal{P}$-ontic), then his ex ante set of priors has a special shape (it is $\mathcal{P}$-rectangular). It can be understood in the terms introduced in Section IX.3.3, namely, as concerning the relationship between the decision maker's current beliefs about the state of the world (the ex ante set of priors) and his beliefs about his possible future states of belief (the epistemic contingencies he envisages). By contrast with Proposition IX.2, which was relatively weak, Proposition IX.4 identifies a particularly strong relationship in a special case. First of all, it implies that whenever the agent's envisaged future beliefs have a particular form – the epistemic contingencies are $\mathcal{P}$-ontic – the aggregator by which the current beliefs are related to future ones is a maxmin EU one (of the sort discussed in Section IX.3). Furthermore, and more importantly, it tells us that this particular form of the envisaged ex post beliefs implies that the ex ante beliefs must also have a specific form: that is, they must be $\mathcal{P}$-rectangular. [10] Note that this form can be relatively straightforwardly verified on inspection of the ex ante set of priors. (It suffices to check that the set $\{p(\bullet|A_i) \mid p \in \mathcal{C},\ p(A_i) = x\}$ is the same for all x.) This renders the Proposition more interesting, because it makes the contrapositive of the implication meaningful.

Indeed, taking the contrapositive, Proposition IX.4 implies that (assuming Dynamic Consistency) if the decision maker's ex ante set of priors is *not* $\mathcal{P}$-rectangular – it does not have the particular shape – then his set of epistemic contingencies is not $\mathcal{P}$-ontic – he does *not* think that he will necessarily learn exactly one event in the partition $\mathcal{P}$ and only that. In other words, it allows the theorist to draw a conclusion about the future beliefs the decision maker anticipates – and hence the subjective decision tree he is using – on the basis of a property of his current beliefs about the state of the world.

To see the consequences for the application of non-expected utility models in dynamic problems, consider the argument with which this section began. It is straightforward to check that the set of priors given above is not $\{Y, Y^c\}$-rectangular. As Proposition IX.4 makes clear, it does not follow that this decision maker is not Dynamically Consistent: under the assumption of Dynamic Consistency, it simply follows that he does not anticipate learning Y and learning Y^c as the only possible future states of belief. In other words, this ex ante set of priors implies that his subjective decision tree is not the objective decision tree given in Figure IX.2. Hence when he is told that he will be faced with the decision tree given in Figure IX.2, he is learning something: namely, something about the future contingencies that need to be planned for. Given that, as has been discussed in the previous section, DC implies that the decision maker's beliefs concerning his future beliefs place constraints on his current beliefs, it should not be surprising that when he updates his beliefs about his future beliefs – to incorporate the information that the only two fu-

10. Note that the use of the maxmin EU aggregator alone does not imply the special rectangular shape of the ex ante set of priors: that also requires that the supports of the ex post sets of priors are disjoint.

ture contingencies are learning Y and learning Y^c – this has consequences for his current beliefs. Indeed, Proposition IX.4 implies that when (or if) he updates his beliefs about his future beliefs to 'adopt' the decision tree in Figure IX.2, then he must also update his current beliefs so that they are $\{Y, Y^c\}$-rectangular. Since the subjective tree then coincides with an objective tree, the standard form of dynamic consistency will thus be satisfied, with respect to the ex ante beliefs the decision maker has *after* having updated on the information about the decision tree that he is facing. The decision maker has no trouble satisfying Dynamic Consistency in this example: it is always satisfied with respect to the subjective decision tree he thinks he is faced with and the ex ante beliefs he has when he thinks he is faced with that tree. The purported problem posed by the example is the result of the inappropriate comparison of his ex ante set of priors before learning that he is faced with the tree in Figure IX.2 and his subjective decision tree after having learnt that he will be faced with the tree in Figure IX.2.

We can thus draw two morals for the use of the maxmin expected utility model in dynamic problems. [11] First of all, Dynamic Consistency can be assumed to be universally satisfied without problem: indeed, given the reasonable intuition supporting it, it is a very natural assumption. Secondly, DC implies a particular relationship between the decision maker's current set of priors and the decision tree he considers himself to be faced with. Any theorist modelling decision makers must of course respect this relationship, if he is to avoid making inconsistent modelling assumptions.

The concrete advice for applications to dynamic problems is simple: when applying the maxmin expected utility model, the ex ante set of priors should be rectangular with respect to the partition formed by the nodes in the tree. It is known, since the seminal work of Epstein and Schneider (2003), that this is the only case in which the decision maker will be dynamically consistent in the classic sense, and hence where standard techniques such as folding back can be used. However, it is sometimes suggested that the restriction to rectangular sets of priors is at best a partial solution, for the problem of dynamic inconsistency remains when the set of priors is not rectangular with respect to an event of interest (see, for example, Hanany and Klibanoff (2007, p282) and Al Najjar et al. (2009, §3.4)). [12] Proposition IX.4 provides a reply to this worry, in the form of a principled justification of

11. These points hold *mutatis mutandis* for other non-expected utility models.

12. Epstein and Schneider's own conclusion (p14) is that 'in *some settings*, ambiguity may render dynamic consistency problematic.'

the restriction to rectangular sets. If the set of priors is not of the appropriate rectangular form, this does not imply that the decision maker is dynamically inconsistent, in the refined sense given in Section IX.3. Rather, it implies that the decision maker is not using the same decision tree as the theorist. So the restriction to appropriately rectangular sets amounts to an assumption that the decision maker does indeed know the decision tree he is faced with, and that this is the same as the one the theorist is using. As such, there is nothing arbitrary or unnecessarily limitative about it; indeed, without such assumptions about the way decision makers conceive the problems they are faced with, economic modelling can hardly get off the ground!

We thus conclude that the proposed notion of dynamic consistency, and the subjective framework for thinking about such issues, has no trouble coping with the purported difficulties for non-expected utility models in dynamic choice problems. Quite to the contrary, it provides a reasoned defense of a certain approach to applying non-expected utility models to such problems, which has already been proposed in the literature. Non-expected utility models can be used in dynamic problems, and even using folding back reasoning: this is a consequence of the fact that there is no obstacle to non-expected utility decision makers being dynamically consistent, with respect to decision tree they consider themselves to be using. However, when using these models, one must respect the consequences of one's assumptions about what the decision maker takes the decision tree he is facing to be: the lesson of Proposition IX.4 is that many purported violations of dynamic consistency are just modelling errors on the part of the theorist.

Remark IX.5. Note that singleton sets of priors are the only sets that are $\mathcal{P}$-rectangular for every partition $\mathcal{P}$. Consequently, the only ex ante preferences such that, for every partition $\mathcal{P}$, there exists a $\mathcal{P}$-ontic set of epistemic contingencies with respect to which DC is satisfied, are expected utility preferences. This is a version of the standard result concerning SDC and expected utility, recovered in our framework. In the light of the preceding considerations, this means that expected utility preferences are the only ones from which one cannot immediately draw the sorts of implications just described about the decision tree that the decision maker is using. The potential for the sorts of modelling errors discussed above is thus reduced with expected utility preferences.

IX.5 Value of Information

In this section, we consider another important economic issue raised in discussions of dynamic consistency, the attitude towards information. A commonly noted purported consequence of non-expected utility is that sophisticated decision makers may prefer not to obtain free information: they are *information averse* (Wakker, 1988; Al Najjar et al., 2009). This of course is in stark contrast with basic pre-theoretical intuition, as well as with the standard Bayesian expected utility model, and is often taken to be an argument, if not one of the strongest arguments, against non-expected utility theory. As a further illustration of the power of the approach proposed here, we analyse the issue of value of information in the framework introduced in the previous sections.

IX.5.1 Information aversion: debunking the standard argument

The standard argument that non-expected utility decision makers are information averse can be formulated on a sequential extension of the dynamic Ellsberg example given in Section IX.1.2. Consider a decision maker with standard Ellsberg preferences – he prefers g_1 to f_1 at node A_1 in Figure IX.2(a), but prefers f_{RB} to g_{RB} at node B_1 in Figure IX.2(b) – and suppose that he is given the choice between facing these two decision trees. [13] That is, he is faced with the decision tree in Figure IX.4. Since he has the Ellsberg preferences just specified, he knows that if he reaches decision node A he will choose g_1 and that if he reaches decision node B he will choose f_{RB}. Reasoning by backwards induction, [14] the decision maker at node C knows that if he takes the upper branch NL of the tree, he will essentially end up with g_1, and if he takes the lower branch L he will end up with f_1 (given that the decision maker does not know the resolution of the uncertainty at the chance node $*$, the choice of f_{RB} at node B essentially boils down to a choice of f_1 from the point of view of a decision maker at C). Hence, since he prefers g_1 to f_1, he chooses NL at C. However, since the choice at C is essentially that between learning whether the ball is yellow before deciding to bet (option L) or not learning whether it is yellow (option NL), by choosing

13. In this discussion, notations and numbering are taken from Section IX.1.2.

14. Decision makers who reason in this way in situations of dynamic inconsistency are sometimes said to adopt 'the strategy of consistent planning', or to be 'sophisticated'; see Strotz (1955); Karni and Safra (1990); McClennen (1990); Siniscalchi (2011).

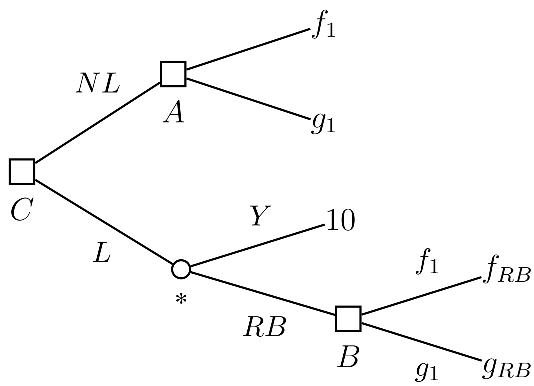

Figure IX.4 – Information Aversion

NL, the decision maker betrays a preference for not obtaining free information. This is the criticized information aversion.

In Section IX.4, we noted that the assumption of Dynamic Consistency implies a strong modelling constraint on the theorist, which undermines the standard practice of taking a decision maker with a given set of priors, subjecting him to a decision tree, and deducing that he is dynamically inconsistent. An essential insight is that, when the decision maker learns that he is faced with a particular decision tree, this may affect his current set of beliefs. The same general point applies in the sequential situation described here. Let $\mathcal{C}$ be the set of priors that decision maker has after having chosen not to learn – that is, at node A. As specified above, this set of priors yields the standard Ellsberg preferences, and so is not rectangular with respect to the partition $\{Y, Y^c\}$. However, if the decision maker choses to learn, then *even before learning whether the ball is yellow or not* – that is, even at the chance node $*$ – his set of priors will not be $\mathcal{C}$. This is a consequence of Proposition IX.4: at node $*$ he knows that he is facing an objective tree which will determine which element of the partition $\{Y, Y^c\}$ occurs, and so his set of priors must be rectangular with respect to this partition (assuming Dynamic Consistency). Since $\mathcal{C}$ is not rectangular with respect to this partition, it follows that his set of priors at node $*$ is not $\mathcal{C}$. In other words, his choice of whether to learn or not at node C affects the beliefs he has immediately after his decision, but before any (other) learning has taken place.

This is a natural consequence of the discussion in the previous sections. If aspects of the decision maker's future beliefs are reflected in his current beliefs, then choices he makes that determine the future beliefs he envisages as possibly having – such as the choice of whether to obtain information or not – impact on his current beliefs. This impact occurs even before he acquires any of the promised information: it is, if you will, the information that he will acquire some information that induces this change in current beliefs, rather than the information that will be acquired itself. And this 'second-order' information comes directly from the choice he has made. As such, there is a strong analogy between the choice of whether to learn and moral hazard problems. The decision maker above is in a similar situation to an Agent whose chances of a particular reward under a given contract depend on the effort he chooses to expend. In both cases, the decision makers' beliefs (about the colour of the ball, or about whether he will get the reward respectively) are directly determined by their choices.

This analogy suggests a diagnosis of the argument sketched above purporting to establish information aversion. The argument considers that the choice between NL and L at node C boils down to the choice between g_1 and f_1. It assumes that the decision maker has the Ellsberg set of priors $\mathcal{C}$ at node C, and evaluates the choice at that node on the basis of the preferences over f_1 and g_1 corresponding to these priors. However, this amounts to the assumption that the decision maker's set of priors immediately after his choice at C are independent of the choice made, and we have just seen that this is false. Since the choice at node C alters his beliefs and hence his evaluations of acts, the decision maker should use the interim beliefs (or if you wish, the beliefs conditional upon the choice to learn or not) to evaluate the options NL and L at node C. This is the standard approach in sequential choice situations of this sort. Consider an Agent who has to choose a level of effort, and then will have the choice of which contract to sign, before observing the outcome of his work (eg. the profits); assuming two levels of effort and two contracts, the decision tree for this problem is given in Figure IX.5. Under the standard approach, one would not calculate ex ante (ie. before knowing the choice of effort) whether the variable or fixed contract is preferable, and then choose whatever effort level leads to that choice of contract at the Contract nodes. Such reasoning, which is the analogue of the above argument for information aversion, obviously errs by wrongly assuming that the optimal choice of contract is independent of the effort exerted.[15] Indeed, the standard approach is to evaluate each

15. Another problem with the reasoning is that it relies on an ex ante probability distribution, which depends on the choice of effort and so requires the Agent to estimate the probability of making a particular choice in the decision he is currently deliberating upon. Similarly in the information case, properly determining the decision maker's set of priors at node C in the tree in Figure IX.4 would require his beliefs about

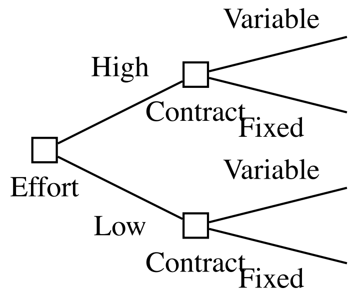

Figure IX.5 – Agent's choice

of the available strategies – (High effort, variable contract), (High effort, fixed contract) and so on – and these evaluations would use the probabilities (beliefs) of outcomes conditional on the choice of effort made. The analogue of this sort of calculation in the case of the choice of obtaining information does not immediately imply that non-expected utility decision makers are information averse.

IX.5.2 Value of envisaged information

What, then, can be said about the value of information under the current proposal? Since the proposal incorporates the beliefs the decision maker envisages himself as possibly having, and hence implicitly the information he expects to receive, the question can be naturally divided into two. On the one hand, what is the value of the information that he expects to receive? On the other hand, what is the value of information that the decision maker is offered but had not necessarily expected to receive? For expositional purposes, it is convenient to treat the former question first.

The issue is essentially whether a decision maker would accept to delay his decision, on the (subjective) decision tree he envisages, until he arrives at the future nodes. It is sometimes suggested that this is immediately implied by dynamic consistency (see for example Ghirardato (2002)); however, such an interpretation involves identifying contingent plans with pairs of acts, and hence it does not easily transfer over to the case of subjective trees. Nevertheless it can be shown, by properly defining the value of information in a manner analogue to the standard Bayesian literature, that it is indeed non-negative.

whether he will choose to learn or not at that node. Doubts may be raised as to whether such beliefs can be formed or be economically meaningful.

The value of information over envisaged contingencies involves the comparison between the value of choosing now and the value of waiting for the realisation of the contingency before choosing.[16] Unlike the standard literature, the ex post value of the information and the ex ante value of the option of not receiving the information are calculated using the maxmin expected utility rule with the ex post and ex ante sets of priors. However, to calculate the ex ante value of obtaining the information, an 'aggregator' connecting ex post and ex ante preferences is required. (In the standard Bayesian case, the aggregator is the expected utility rule.) We thus assume that there is such an aggregator $\phi : \Re^I \to \Re$: i.e. that there exists ϕ such that, for all $f \in \mathcal{A}$, $\min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s) = \phi((\min_{p \in \mathcal{C}_i} \sum_{s \in S} u(f(s))p(s))_{i \in I})$. By Proposition IX.2, such an aggregator exists whenever the decision maker satisfies DC. Under this assumption, the following definition of non-negative value of information is correctly formulated.

Definition IX.2. Let $\mathcal{C}$ and $\{\mathcal{C}_i\}_{i \in I}$, for $i \in I$ represent ex ante and ex post sets of priors and let $\phi : \Re^I \to \Re$ be the aggregator. The *value of envisaged information for a menu* $A \in \wp(\mathcal{A})$ *is non-negative* if and only if:

$$\phi\left(\left(\max_{f \in A} \min_{p \in \mathcal{C}_i} \sum_{s \in S} u(f(s))p(s)\right)_{i \in I}\right) \geqslant \max_{f \in A} \min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s) \qquad \text{(IX.2)}$$

The *value of envisaged information is always non-negative* if it is non-negative for every $A \in \wp(\mathcal{A})$.

This notion of non-negative value of information is evidently analogous to the standard notion, the difference being that it does not involve the assumption of expected utility decision makers, and is specific to the information that the decision maker envisages acquiring ex ante – that is, to his subjective decision tree.

Proposition IX.5. *If the decision maker satisfies DC, then the value of envisaged information is always non-negative.*

For the information the decision maker envisages receiving, the refined notion of Dynamic Consistency implies that it has non-negative value. The standard point concerning

16. We are following a standard approach to information value in, for example, Marschak and Miyasawa (1968); Marschak and Radner (1972); Hilton (1981).

SDC and objective trees thus continues to apply for DC and the decision maker's subjective tree: he weakly prefers to wait for the information before deciding. Here there is no possibility of information aversion.

IX.5.3 Value of offered information

The situation is more subtle in cases where the decision maker had not expected to learn (just) the information on offer, and so learns something about his potential future beliefs – or perhaps more accurately, about the fact that they will be modified – on choosing to obtain the information. As explained above, such cases are analogous to moral hazard situations, and should be analyzed accordingly. In particular, the value of information involves the comparison of the value of choosing just after having rejected the offer of information, and the value of accepting the offer and choosing after having received the information. As is clear from the discussion in Section IX.5.1, this is not the standard way of reasoning about attitude to information in the literature on non-expected utility theory.

For the analysis, consider an offer of acquisition of information $\mathcal{I}$. Let $\mathcal{C}$ be the decision maker's set of priors immediately after having chosen not to obtain the information, and let $\mathcal{C}_{(\mathcal{I})}$ be his set of priors immediately after having chosen to obtain the information, but before actually receiving it. Being slightly slack with terminology, we shall continue to call these the ex ante sets of priors in the two cases. Let $\{\mathcal{C}_i\}_{i\in I}$, $\phi : \Re^I \to \Re$ and $\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}$, $\phi_{(\mathcal{I})} : \Re^K \to \Re$ be the ex post sets of priors and aggregators just after having chosen not to turn down and acquire the information respectively. [17]

Remark IX.6. Note that this framework is broad enough to cover standard representations of information structures or systems. One could consider a set of signals, and understand each $\mathcal{C}_{(\mathcal{I})k}$ to be a posterior set of beliefs after the reception of some signal, with $\phi_{(\mathcal{I})}$ reflecting the decision maker's prior beliefs as to which signal will be received. As such, the framework essentially corresponds to a standard representation of information structures (for example Gollier (2004, Ch 24)), the only difference being that the beliefs and aggregator are not necessarily expected utility. Similarly, another standard representation of information structures as partitions (eg. Marschak and Miyasawa (1968); LaValle (1968)) can

17. It is natural to assume that $\mathcal{C}$, $\{\mathcal{C}_i\}_{i\in I}$ and ϕ are independent of the information $\mathcal{I}$ offered, insofar as they can be thought of as representing the decision maker's state even before he envisaged the possibility of being offered the information. Nothing in the analysis below depends on this assumption.

be incorporated, by adding assumptions that the support of each $\mathcal{C}_{(\mathcal{I})k}$ belongs to exactly one element of the given partition. [18]

Moreover, this setup is silent on the question of the relationship between the beliefs in the case where the information is rejected and those in the case where it is accepted. The development of update rules that determine how the decision maker's envisaged future beliefs are revised upon learning that he will receive some hitherto unexpected information is a topic for future research.

The standard question of value of information involves comparison of the value of deciding immediately after having chosen not to receive the information – so the relevant set of priors is $\mathcal{C}$ – with the anticipated value of deciding after having received the information – so the relevant elements are the sets of priors $\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}$ and the opinion as to which will be realized, given by $\phi_{(\mathcal{I})}$. Hence the following definition of the non-negative value of offered information.

Definition IX.3. Let $\mathcal{C}$, $\{\mathcal{C}_i\}_{i\in I}$, ϕ, $\mathcal{C}_{(\mathcal{I})}$, $\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}$, $\phi_{(\mathcal{I})}$ be the ex ante beliefs, envisaged ex post beliefs and aggregators in the case where the decision maker chooses to reject the information $\mathcal{I}$ or to accept it respectively, for some offered information $\mathcal{I}$. The *value of offered information $\mathcal{I}$ for a menu $A \in \wp(\mathcal{A})$ is non-negative* if and only if:

$$\text{(IX.3)} \qquad \phi_{(\mathcal{I})}\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u(f(s))p(s)\right)_{k\in K}\right) \geqslant \max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s)$$

The *value of offered information $\mathcal{I}$ is always non-negative* if it is non-negative for every menu $A \in \wp(\mathcal{A})$.

Envisaged information is recovered as the special case where the decision maker is offered information which he already expected to learn (ie. where $\mathcal{C} = \mathcal{C}_{(\mathcal{I})}$, $\{\mathcal{C}_i\}_{i\in I} = \{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}$, $\phi = \phi_{(\mathcal{I})}$). In other cases, there are clear non-trivial necessary and sufficient conditions for the value of offered information to be non-negative.

18. Note that there is not necessarily a one-to-one correspondence between the 'objective' signals or cells of the partition and the ex post sets of priors, because the decision maker may have expected to learn other information over and above the information offered. For example, if he expects to learn which of the events B or B^c hold, and is offered the information of which of A and A^c hold, then there may be up to four relevant ex post sets of priors – after learning A and B, after learning A and B^c, and so on – though the size of the information partition offered is two.

Proposition IX.6. *Suppose that the decision maker satisfies DC. Then the value of information $\mathcal{I}$ is always non-negative iff $\mathcal{C}_{(\mathcal{I})} \subseteq \mathcal{C}$.*

This proposition provides the final piece in the analysis of the allegation that non-expected utility decision makers necessarily exhibit information aversion. It tells us that the value of information for any maxmin expected utility decision maker will be non-negative as long as his set of priors immediately after having chosen to learn the information is contained in the set of priors he has immediately after having chosen not to acquire the information. Since the standard arguments, such as the one discussed in Section IX.5.1, involve objective trees, the information offered comes in the form of a partition and this adds another constraint: namely, that the set of priors immediately after accepting the information are $\mathcal{P}$-rectangular.[19] These two constraints are simultaneously satisfiable: for one, the singleton sets containing priors in $\mathcal{C}$ satisfy both conditions. Hence, for any decision maker who has a set of priors $\mathcal{C}$ if he decides not to accept the information, there exists a set of priors such that, if he holds this set of priors after having chosen to acquire the information, he will have non-negative value of information. In other words, a non-expected utility decision maker can always exhibit weak preference for information with respect to any information partition offered. Once the fact that the decision maker may change his beliefs after choosing to acquire information is properly incorporated into the calculation of the value of that information, it is clear that claim that non-expected utility necessarily implies information aversion is not only unfounded: it is false.

IX.5.4 Negative value of information and opportunity costs

One might nevertheless remain puzzled by the fact that it is possible, if the constraint in Proposition IX.6 is not met, for a decision maker to give a negative value to some information (for some menu). Is the mere possibility of turning down free information not itself an unacceptable consequence of non-expected utility? There is certainly a strong intuition that one should never have a strict preference for refusing free information; but it implicitly rests upon the assumption that the information acquisition does not compromise the reception of information one had otherwise expected to obtain. Suppose that James,

19. Since, as noted in footnote 18, there may be more than one ex post preference that is non-null on a given cell of the partition, Proposition IX.4 does not strictly speaking apply. However, it is straightforward to extend it to cover this case.

a manager of a firm, has a spy in a competitor's board of directors, and expects to learn from the spy the competitor's strategy in a market B. He is offered the possibility to learn 'for free' the competitor's strategy in another market A. If this information can be acquired without consequence for the prospects of gaining information about market B, then it is evidently of non-negative value. Suppose, however, that acquiring the information about the strategy in market A involves blowing the spy's cover, so he will no longer be able to tell James anything about market B. In this case, it may well be reasonable to refuse the offer of information about market A, because acquiring this information jeopardizes other information that James had expected to receive. Of course, in this sort of case, the choice is essentially between two different pieces of information: the information the decision maker expected to get before the offer and the information he will get if he accepts the offer. If the value of the former piece of information is greater than that of the latter – if the opportunity cost of the offered information outweighs its value – then it is perfectly reasonable to turn down the offer of 'free' information. The condition in Proposition IX.6, we suggest, can be understood as reflecting this caveat.

A clear way to see this point is by considering what happens in situations where the assumptions underlying the standard intuition in favour of non-negative value of information are satisfied: that is, in cases where choosing to acquire the offered information does not imply forgoing information that one already expected to obtain. Note firstly that, as mentioned above (see Remark IX.6), the ex post priors and aggregators in the case where the information is declined ($\{(\mathcal{C}_i\}_{i\in I}, \phi)$) and the ex post priors and aggregators in the case where the information is accepted ($\{(\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}, \phi_{(\mathcal{I})})$) can be thought of as information structures, and the decision of whether to accept the information is effectively a choice between these two information structures. If no previously expected information is compromised on accepting the offer of information, then the latter information structure is more informative than the former one. According to a standard definition, informativeness can be defined in terms of decision makers' attitudes to the information: one information structure is more informative than another if every decision maker (no matter his utility function) would prefer to learn according to the former information structure (Marschak and Miyasawa, 1968; Gollier, 2004).[20] This definition can be adapted to the current framework as

20. Naturally, there are a battery of important results showing that this definition is equivalent to other possible definitions of informativeness under expected utility (Blackwell, 1953; Marschak and Miyasawa, 1968; Gollier, 2004). We are not aware of similar results for non-expected utility models; indeed, it is not

follows.

Definition IX.4. Let $\mathcal{C}$, $\{\mathcal{C}_i\}_{i\in I}$, ϕ, $\mathcal{C}_{(\mathcal{I})}$, $\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}$, $\phi_{(\mathcal{I})}$ be the ex ante beliefs, envisaged ex post beliefs and aggregators in the case where the decision maker chooses to reject the information $\mathcal{I}$ or to accept it respectively. $(\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}, \phi_{(\mathcal{I})})$ is *at least as informative as* $(\{\mathcal{C}_i\}_{i\in I}, \phi)$ if and only if, for every affine utility function $u' : \Delta(X) \to \Re$, and every menu $A \in \wp(\mathcal{A})$:

(IX.4)
$$\phi_{(\mathcal{I})}\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u'(f(s))p(s)\right)_{k\in K}\right) \geqslant \phi\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_i}\sum_{s\in S}u'(f(s))p(s)\right)_{i\in I}\right)$$

It follows directly from this definition and Proposition IX.5 that, whenever the information structure one faces after accepting the information offered is at least as informative as the information structure faced when declining the offer, then the value of the offered information is non-negative.

Proposition IX.7. *Suppose that the decision maker satisfies DC. Then the value of information $\mathcal{I}$ is always non-negative whenever $(\{\mathcal{C}_{(\mathcal{I})k}\}_{k\in K}, \phi_{(\mathcal{I})})$ is at least as informative as $(\{\mathcal{C}_i\}_{i\in I}, \phi)$.*

Hence, in accordance with the standard intuition, non-expected utility decision makers will not turn down free information when it does not compromise information they had previously expected to receive.

We conclude that, even for non-expected utility decision makers, their attitude to information behaves exactly 'as it should do'. On the one hand, the fact that they are non-expected utility does *not* imply that they will *necessarily* ever have negative value of information: Proposition IX.6 implies that it is always possible for their beliefs after having chosen to acquire the information to be such that the value of the information is non-negative. On the other hand, it is *possible* for them to exhibit negative value of information – the condition in Proposition IX.6 is non-trivial. However, when this happens, it does so for good reason. In particular, Proposition IX.7 tells us that, whenever the information offered does not compromise information that was expected to be acquired – whenever it is a simple addition of information, with no subtraction, so to speak – the value they assign to

even evident what the analogues for some of the other proposed definitions would be for such models.

the information offered is non-negative. If the value of the information offered is negative, then this must be a case where the decision maker forgoes information he had expected to receive on accepting the information offered. In these cases, aversion towards the information offered is entirely acceptable: if the information lost by accepting the offer is worth more than the information acquired, then it is entirely reasonable to decline the offer.

IX.6 Foundations

So far we have assumed a set of epistemic contingencies as given; however, they are evidently not directly observable. The approach depends on the assumption that the decision maker does envisage such epistemic contingencies; moreover, it is important, to ascertain whether the refined notion of Dynamic Consistency is respected or not, to determine which epistemic contingencies the decision maker envisages. The aim of this section is to present behavioral foundations for the notion of epistemic contingency in the context of the maxmin EU model. To this end, we shall introduce a choice-theoretic situation and propose axioms that are necessary and sufficient for a representation of choices involving envisaged sets of future priors $\{\mathcal{C}_i\}_{i\in I}$. The result in this section should relieve worries that the proposed reconceptualization of dynamic choice via the introduction of the notion of epistemic contingency relies on notions with no independent behavioral meaning. The notions have in-principle observable consequences (beyond the application to the question of dynamic consistency), and the set of epistemic contingencies envisaged by the decision maker can in principle be gleaned from choice behavior.

IX.6.1 Setup and representation

We continue to work in the setup set out in Section IX.2.1. Rather than assuming (future) sets of priors or preference relations as primitives, we work now with choice functions, where a choice function is a function $c : \wp(\mathcal{A}) \to \wp(\mathcal{A})$ such that, for any $A \in \wp(\mathcal{A})$, $c(A) \subseteq A$.[21] The choice function c has the following interpretation, associated with the time line given in Figure IX.6. The decision maker knows that he will have to choose an act from a menu A at an ex post stage (before the realisation of the state of the world, but

21. This is sometimes called a choice correspondence in the literature.

perhaps after receiving information), and has the opportunity in an ex ante stage of restricting the options left open to a subset of A, from which he will make his ex post choice. For an example of such a choice situation, consider a committee deciding on the allocation of a building contract or of a university post: they may in the first instance rule out some of the candidates, producing a shortlist, from which they will later choose the winner. For any $A \in \wp(\mathcal{A})$, c yields the set of elements that the decision maker now wishes to keep as open alternatives for his future choice; the elements not in $c(A)$ are those which the decision maker is willing to rule out now.

We shall consider the following representation:

$$\text{(IX.5)} \qquad c(A) = \{f \in A \mid f \in \arg\max_{g \in A} \min_{p \in \mathcal{C}'} \sum_{s \in S} u(g(s))p(s) \text{ for some } \mathcal{C}' \in \mathcal{P}\}$$

where u is a continuous affine utility function on $\Delta(X)$ and $\mathcal{P}$ is a set of convex, closed subsets of $\Delta(\Sigma)$. The sets of priors in $\mathcal{P}$ are interpreted as the future beliefs that the decision maker anticipates himself as possibly having at the moment when he will be faced with his final choice; they can be thought of as the epistemic contingencies he envisages. (IX.5) represents a decision maker who envisages that he will hold one of the sets in $\mathcal{P}$ as his set of priors at the specified future moment, and that he will form preferences with that set according to the maxmin EU rule. He retains as an open option any act which is optimal according to this rule with any of the sets of priors in $\mathcal{P}$. To see this more clearly, note that this representation can be reformulated in terms of the ex post preferences in the different epistemic contingencies, as follows. Define the set $\mathcal{R}$ of preference relations on $\mathcal{A}$ as follows: $\preceq_i \in \mathcal{R}$ if and only if there exists $\mathcal{C}_i \in \mathcal{P}$ such that $\preceq_i$ is represented according to (IX.1) by $\mathcal{C}_i$ and u. $\mathcal{R}$ contains the decision maker's anticipated ex post preferences, formed on the basis of the ex post beliefs according to the maxmin EU rule. Representation (IX.5) is equivalent to the following representation of ex ante choice in terms of a set $\mathcal{R}$ of preference relations each of which is represented according to (IX.1):

$$\text{(IX.6)} \qquad c(A) = \{f \in A \mid \exists \preceq_i \in \mathcal{R},\ \forall g \in A,\ f \succeq_i g\}$$

In other words, the decision maker keeps open any act which is optimal in the menu according to at least one of his anticipated ex post preferences, and rules out any act which

Figure IX.6 – Time line

ex ante stage	ex post stage	realisation and payoff
choose shortlist $B \subseteq A$	choose act $f \in B$	

is not optimal under any of the ex post preferences. The choice rule (IX.6) yielding a choice function given a set of preferences has been studied in the choice-theoretic literature, where it has been called *pseudo-rationalisability* (Moulin, 1985).

Remark IX.7. Note that representation (IX.5) could alternatively be formulated in terms of preferences over menus. To illustrate the relationship, define the preference relation $\dot{\prec}$ on $\wp(\mathcal{A})$ as follows: for all $A, B \in \wp(\mathcal{A})$,

$$A \dot{\prec} B \quad \text{iff} \quad c(A \cup B) \cap A = \varnothing \tag{IX.7}$$

It is straightforward to check that this preference over menus is represented by the u and $\mathcal{P}$ featuring in representation (IX.5) as follows: for all $A, B \in \wp(\mathcal{A})$, $A \dot{\prec} B$ iff

$$\max_{g \in A} \min_{p \in \mathcal{C}'} \sum_{s \in S} u(g(s))p(s) < \max_{g \in B} \min_{p \in \mathcal{C}'} \sum_{s \in S} u(g(s))p(s) \ \text{ for all } \mathcal{C}' \in \mathcal{P} \tag{IX.8}$$

The relation $\dot{\prec}$ is a preference relation over menus in the style of Kreps (1979); Dekel et al. (2001), but, unlike most preference relations studied in this literature, it will not satisfy the standard conditions on preferences: in particular, the weak version of the relation will not be complete. This is due to the representation. It is standard in the aforementioned literature to assume that ex ante preferences are formed using an 'aggregator' that ensures completeness; by contrast, representation (IX.8) combined with definition (IX.7) involves a 'neutral' aggregator that only orders menus in cases of strict dominance. This difference is related to the goal of the exercise in this section. Whereas most of the representation theorems in the preference for flexibility literature deliver both ex post preferences (or states of a subjective state space) and ex ante preferences that are aggregated in a particular way with respect to the ex post ones, here the aim is to give foundations solely for ex post sets of priors or preferences (the epistemic contingencies), without making any assumptions about the ex ante preferences. Indeed, the issue of the relationship between ex post and ex ante preferences is essentially a question about dynamic consistency, and this should evidently be kept separate from the behavioral foundation of the notion of epistemic contingency.

IX.6.2 Axioms and result

Consider the following axioms on the choice function c.

Axiom IX-A1 (Chernoff)**.** For all $A, B \in \wp(\mathcal{A})$, $f \in \mathcal{A}$, if $A \subseteq B$ and $f \in c(B)$, then $f \in c(A)$.

Axiom IX-A2 (Aizerman)**.** For all $A, B \in \wp(\mathcal{A})$, if $c(B) \subseteq A \subseteq B$, then $c(A) \subseteq c(B)$.

Axiom IX-A3 (Non-degeneracy)**.** There exist $d, e \in \Delta(X)$ such that $d \in c(\{d, e\})$ and $e \notin c(\{d, e\})$.

Axiom IX-A4 (Fixed utilities)**.** For all $d, e \in \Delta(X)$ and $A, B \in \wp(\Delta(X))$ with $A \subseteq B$, if $d, e \in c(A)$, then $d \in c(B)$ if and only if $e \in c(B)$.

Axiom IX-A5 (Set C-Independence)**.** For all $A \in \wp(\mathcal{A})$, $d \in \Delta(X)$, and for all $\alpha \in (0, 1)$, $c(A_\alpha d) = c(A)_\alpha d$.

Axiom IX-A6 (Union C-Independence)**.** For all $A \in \wp(\mathcal{A})$, $\alpha \in (0, 1)$ and $d \in \Delta(X)$ with $d \in c(A)$, $c(A) \subseteq c(A \cup A_\alpha d)$.

Axiom IX-A7 (Monotonicity)**.** For all $A, B \in \wp(\mathcal{A})$ with $A \subseteq B$, if, for each $g \in B$, there exists $f \in A$ with $f(s) \in c(\{f(s), g(s)\})$ for all $s \in S$, then $c(A) \subseteq c(B)$. Moreover, for every $g \in B$, if there exists $f \in B$ with $g(s) \notin c(\{f(s), g(s)\})$ for all $s \in S$, then $g \notin c(B)$.

Axiom IX-A8 (Uncertainty aversion)**.** For all $A, B \in \wp(\mathcal{A})$ with $B \subseteq c(A)$, $f \in c(A \cup \{f\})$ for all $f \in conv(B)$ whenever there exist $d \in \Delta(X)$ and $g \in B$ such that: i. $f_\alpha d \in B$ for all $\alpha \in [0, 1]$ and all $f \in B$; and ii. $g_\alpha d \notin c(A \cup \{f_\beta e\})$ for all $f \in B$, all $e \in \Delta(X)$ with $e \succ d$ and all $\alpha, \beta \in (0, 1)$.

Axiom IX-A9 (Continuity)**.** For all sequences of menus $(A_n)_{n \in N}$ and $A \in \wp(\mathcal{A})$ with $A_n \to A$ and all sequences of acts $(f_n)_{n \in N}$ with $f_n \in A_n$ for each $n \in N$, if $f_n \to f$, then $f \in c(A)$.

Chernoff (IX-A1) and Aizerman (IX-A2) are standard axioms in the choice-theoretical literature. The former is just Sen's axiom α, and requires no further comment; the latter says that removing an alternative that does not belong to the choice set cannot lead to new alternatives being chosen. The conjunction of the two is weaker than the Weak Axiom of Revealed Preference, and equivalent to the pseudo-rationalisability of c (representation (IX.6)); the reader is referred to Moulin (1985), from whom we borrow the nomenclature, for further discussion. Fixed utilities (IX-A4) imposes Sen's axiom β on the restriction of c to menus containing only constant acts. It follows from standard choice theory results (see for example Sen (1971)) that the restriction to constant acts is represented by a single transitive complete preference relation. This axiom translates the assumption that the decision maker's preferences over constant acts are the same in all envisaged future contingencies; he only anticipates differences in beliefs. Non-degeneracy (IX-A3) requires that the choice function is non-trivial on constant acts.

The remaining axioms can be thought of as choice-theoretical analogues of the Gilboa-Schmeidler axioms on preferences (1989). The C-independence axioms (IX-A5 and IX-A6) correspond to Gilboa and Schmeidler's C-independence. Recall that the idea behind their axiom is that mixing with a constant act does not 'change' the preference order. Similarly here, Set C-independence (IX-A5) states that mixing a menu with a constant act does not 'change' the choice set: if the decision maker wanted to keep an act as an open option from a given menu, he would like to keep the mixture of the act as an open alternative from a mixture of the menu. A consequence of the Gilboa-Schmeidler C-independence axiom (in conjunction with other basic preference axioms) is that, for any act f and constant act d, whichever of the acts is weakly preferred between f and d remains weakly preferred over any mix $f_\alpha d$ of the two acts. Union C-independence states the equivalent of this for menus: if the decision maker would keep open an act f and a constant act d from a menu, then adding $f_\alpha d$, or indeed any mixture of d with an element of the menu, does not 'change' his decision to keep f and d open. This translates the idea that if f or d will be possibly chosen in some future contingency when both are available, then $f_\alpha d$ will not be chosen over it. The need for two C-independence axioms is due to the weak choice-theoretical properties that are assumed: if Chernoff (IX-A1) and Aizerman (IX-A2) are replaced by the Weak Axiom of Revealed Preference, it is straightforward to check that Set C-independence implies Union C-independence.

Monotonicity (IX-A7) is essentially the standard monotonicity or statewise dominance axiom formulated for choice sets. It includes both a weak dominance and a strict dominance clause, both of which are fairly intuitive. The first basically says that adding elements to a menu that are weakly dominated by some element already on the menu does not lead one to rule out any of the options that one initially left option. This translates the standard intuition that adding a weakly dominated option should not prevent a previously chosen option from being chosen.[22] The second clause just says that one does not leave strictly dominated acts open: this translates the intuition that strictly dominated options are never chosen. These conditions are in the spirit of the standard monotonicity condition.

Uncertainty Aversion (IX-A8) can be thought of as a weakening of the standard Gilboa-Schmeidler axiom extended to the general menu setting. The standard axiom, formulated on preferences, states that if there is indifference between a pair of acts, then any mixture is weakly preferred to both. A natural extension to the case of general menus is obtained by replacing the pair of indifferent acts by a subset of the menu that is contained in its choice set, and the mixture by any mixture of the elements in the subset. That is, it states that, for all $A, B \in \wp(\mathcal{A})$ with $B \subseteq c(A)$, $f \in c(A \cup \{f\})$ for all $f \in conv(B)$. The axiom IX-A8 is evidently a weakening of this extension, stating that it holds under particular conditions. In fact, the interpretation of the extension requires considering acts in the choice set to be indifferent; however, whilst this is the case under WARP, it is no longer true under the weaker choice-theoretic axioms used here. The conditions in IX-A8 guarantee that there is an ex post preference relation according to which the acts in B are indifferent.[23] So the axiom can be understood as stating that if the decision maker anticipates that he will

22. Note that the proposed interpretation of dominance is vindicated by the Fixed utilities axiom (IX-A4).

23. This can be seen as follows. Whenever a decision maker chooses to leave open the acts f, g, the constant act d, and mixtures $f_\alpha d$ and $g_\alpha d$ from a menu, this means that he envisages contingencies where each of these acts is optimal. Moreover, since, for any ex post maxmin EU preference, $f_\alpha d$ is optimal from a menu containing f and d only if f and d are indifferent (this is a consequence of the Gilboa-Schmeidler axioms), there is an possible contingency in which f, d and $f_\alpha d$ are all optimal. Similarly, there is a possible contingency in which g, d and $g_\alpha d$ are all optimal. Finally, if, for any act dominating $f_\beta d$, adding such an act causes the mixtures $g_\alpha d$ to no longer be left open, this implies that any act dominating $f_\beta d$ is also strictly preferred to g, d and $g_\alpha d$ according to the contingencies where they were all indifferent; so $f_\beta d$, and hence f and d are indifferent to g, d and $g_\alpha d$ under some contingency. Extending this reasoning to more than two acts, the conditions in IX-A8 imply that the decision maker considers it possible that he will be indifferent among the acts in B.

be indifferent between the acts in B, then he anticipates that he will be willing to choose any mixture – since mixtures may hedge the ambiguity in the acts – and so he leaves these mixtures open. As such, IX-A8 captures the hedging intuition in the standard axiom, in the context of the specific interpretation of the choice function employed here.

The final axiom, IX-A9, is the standard upper hemicontinuity property for choice functions. It states that whenever a sequence of menus tends to a menu, then the limit of any sequence of options left open in these menus is left open in the limit menu.

A foundation for the notion of epistemic contingency is given in the following representation theorem.

Theorem IX.1. *The following are equivalent:*

(i) c *satisfies IX-A1–IX-A9;*

(ii) *There exists a nonconstant continuous affine utility function* $u : \Delta(X) \to \Re$ *and a set* $\mathcal{P}$ *of closed convex sets of probability measures on* S *such that:*

$$c(A) = \{f \in A \mid f \in \arg\max_{g \in A} \min_{p \in \mathcal{C}'} \sum_{s \in S} u(g(s))p(s) \text{ for some } \mathcal{C}' \in \mathcal{P}\} \tag{IX.5}$$

Moreover, u *is unique up to positive affine transformation, and there is a unique minimal* $\mathcal{P}$.

This theorem shows that, under certain conditions, a set of sets of priors – or a set of epistemic contingencies – representing the choice function according to (IX.5) exists. Moreover, there is a unique 'canonical' such set, namely the unique minimal set. We conclude that the notion of epistemic contingency does have solid behavioral foundations.

Remark IX.8. Note that the results in this section can be combined with those in Section IX.3 to provide a choice-theoretic axiomatization of a representation in terms of maxmin expected utility ex post preferences and an aggregator that yields complete ex ante preferences. One potential interest for this development would be to translate the discussion in Section IX.3 into purely choice-theoretic terms. One possibility would be to express the Dynamic Consistency condition purely as a relationship between two choice functions – the choice function represented according to (IX.5) and interpreted as discussed above, and a choice function represented according to the Gilboa and Schmeidler (1989) representation (translated into choice-function language) and interpreted in the standard way. With this, one would obtain a choice-theoretic formulation of the notion of dynamic consistency proposed in Section IX.3, as well as a representation result for maxmin EU ex post preferences with an aggregator of the sort given in Proposition IX.2. [24]

Since these formulations do not seem particularly illuminating, we do not present them here.

IX.7 Conclusion

It is commonly held that dynamic consistency, consequentialism and non-neutral attitudes to ambiguity are incompatible. We have argued in this paper that this is not the case, if the dynamic consistency condition is properly formulated. The central idea is that one can only ask a decision maker to be dynamically consistent with respect to the contingencies that he in fact envisages – rather than those imposed by the theorist – and when these contingencies are properly taken into account, the apparent incompatibility is resolved.

To fully defend this rebuttal of the standard argument, we have given choice-theoretic foundations to the central notion of a contingency that the decision maker envisages, on the basis of the choices the decision maker makes to 'leave options open' for subsequent choice at a future time. The representation result can be considered as an independent contribution to the literature on preferences over menus and unforeseen contingencies, and is perhaps the first result of the sort to consider incomplete preferences over menus.

Perhaps even more significant than the defense of this perspective and its ability to resist the arguments against ambiguity models are its conceptual and economic consequences. First of all, it emphasizes that dynamic consistency is fundamentally concerned with the relationship between two sorts of belief: one's current beliefs about the state of the world and one's beliefs about one's possible future beliefs about the state of the world.

Secondly, the perspective provides a principled justification for the use of a restricted family of sets of priors in applications to dynamic choice problems. Applications typically adopt the tacit assumption that the decision maker knows what the decision tree is and that it is the way the theorist describes it. It turns out that Dynamic Consistency, in the refined

24. Along similar lines, replacing the Dynamic Consistency axiom by a choice-theoretic equivalent of the Expert Uncertainty Aversion axiom in Crès et al. (2011), one obtains a representation where the decision maker has maxmin ex post preferences and aggregates them using a maxmin functional.

sense introduced here, implies that this assumption can only hold if the decision maker's ex ante beliefs are of a specific form. A decision maker whose beliefs are not of this form may be perfectly dynamically consistent: he just will not consider himself to be facing the decision tree that the modeler is using. That, of course, is not a problem for the decision maker; it is a problem for the modeler.

Finally, the perspective reveals that the issue of information aversion has been wrongly analyzed in the standard arguments. Non-expected utility decision makers need not be information averse, and will in fact not be whenever the information on offer does not compromise information they had previously expected to receive. They may turn down an offer of 'free' information, but this only occurs when its opportunity cost – in terms of information they had otherwise expected to receive but that is jeopardized by accepting the offer – outweighs its value.

IX.A Incorporating the epistemic contingencies into the state space

It was claimed in Remark IX.2 that the analysis conducted in this paper goes through in other setups; in this Appendix, we illustrate this point by considering a framework where the epistemic contingencies (ie. future states of belief) are explicitly represented in the state space.

We adopt the terminology and assumptions in Section IX.2; recall that S is the 'objective' state space. Let $\mathcal{E}$ be the set of closed convex subsets of $\Delta(\Sigma)$; this is the set of sets of priors over S, and will represent the potential future states of belief (or epistemic contingencies) of the decision maker. The 'extended state space', incorporating the epistemic contingencies, is $\Omega = S \times \mathcal{E}$. The first 'coordinate' s' of a state $(s', \mathcal{C}') \in \Omega$ represents the state of the world that is realized; the second 'coordinate' $\mathcal{C}'$ is the decision maker's future state of belief. Note that this state space is rich enough to represent the arrival of any piece of information: the information is represented by the ex post, or 'conditional', beliefs it induces. Although Ω is rich enough to represent all ex ante uncertainty, both about the state of the world and the ex post state of belief, S is sufficient to represent all payoff-relevant uncertainty. As such, the domain of (ex ante and ex post) preferences is still the set $\mathcal{A}$ defined in Section IX.2, which corresponds naturally to a subset of $\Delta(X)^{\Omega}$. Similar points hold for the sets of priors involved in the representation of preferences, which are defined over S, considered as the partition $\{\{s\} \times \mathcal{E} \mid s \in S\}$ of Ω.

In the setup used above (Section IX.2.1), an ex ante set of priors $\mathcal{C}$ and a set of epistemic contingencies (ex post sets of priors) $\{\mathcal{C}_i\}_{i \in I}$ were assumed. As noted, the former corresponds to a set of priors over the partition $\{\{s\} \times \mathcal{E} \mid s \in S\}$ of Ω. The set of epistemic contingencies corresponds in this setup to a subset of Ω, namely the set $\{(s', \mathcal{C}') \in \Omega \mid \exists i \in I,\ \mathcal{C}' = \mathcal{C}_i\}$. We call this set EC.[25]

Given the lack of a preference relation over $\Delta(X)^{\Omega}$, some remarks are in order about what should count as (the equivalent of a) 'null event' in Ω. On the one hand, states not

25. Whilst we do not assume a preference relation on $\Delta(X)^{\Omega}$ or a set of priors over Ω, the analysis undertaken below continues to hold under such an assumption, as long as the preference relation or set of priors is appropriately consistent with the ex ante set of priors ($\mathcal{C}$) and the set of epistemic contingencies (EC).

belonging to EC can be thought of as 'null' insofar as the decision maker does not consider it possible for them to hold (he does not consider it possible for him to be in the future states of belief coded in these states). Moreover, any state $(s', \mathcal{C}')$ such that s' is null according to the ex ante preference relation $\preceq$ can naturally be thought of 'null'. We will say that $EC \cap NN$ is the set of Ω-*non-null states*, where $NN = \{(s', \mathcal{C}') \in \Omega \mid s' \ \preceq_{ea}$ -non-null$\}$. Any event with non-empty intersect with $EC \cap NN$ will be said to be Ω-*non-null*.

Note that, given the interpretation of the extended state space, it does not make sense to consider learning or conditioning upon certain events of Ω. To characterize the events that could conceivably be learnt, we introduce the following definition.

Definition IX.5. An event $A \subseteq \Omega$ is *learnable* if it satisfies the following two conditions:

1. $\mathcal{C}' = \mathcal{C}''$ for all $(s', \mathcal{C}'), (s'', \mathcal{C}'') \in A$.
2. If $(s', \mathcal{C}') \in A$, then $(s'', \mathcal{C}') \in A$, for every s'' for which $p(s'') > 0$ for some $p \in \mathcal{C}'$.

Moreover, a partition of Ω is *learnable* whenever it consists entirely of learnable events.

Learnable events are those which can conceivably be learnt; this motivates the conditions in the definition.[26] The first condition corresponds to the assumption that, in the ex post stage, the agent has no uncertainty about his ex post beliefs; these are represented by a single set of priors. So the only events in Ω that can be learnt are those which have the same ex post beliefs for every state they contain. The second condition reflects the fact that the set of states of the world that the decision maker considers to be non-null in the ex post stage must respect the set of beliefs he has at this stage. In particular, a state $s' \in S$ cannot be ruled out by the information A (ie. there is no state $(s', \mathcal{C}') \in A$) but nevertheless possibly have non-zero probability according to the ex post beliefs established by A.

Note that, for any learnable event A, the preferences conditional on A, $\preceq_A$, are generated by the set $\mathcal{C}^*$ such that $\mathcal{C}' = \mathcal{C}^*$ for all $(s', \mathcal{C}') \in A$.

We claimed that, with these constraints in mind, this framework yields the same analysis as that carried out above. One way to back up this claim is to show that the standard definition of dynamic consistency (SDC) applied on this framework is equivalent to the refined notion of dynamic consistency proposed in Section IX.3 (DC). Translated into this framework, SDC becomes:

26. Recall that we are working in a simple setup with two time periods. So we are considering the learning of an event between the ex ante and ex post stage.

Standard Dynamic Consistency on Ω. For every $f, g \in \mathcal{A}$ and learnable partition $\{A_j\}_{j\in J}$ of Ω, if $f \preceq_{A_j} g$ for every Ω-non-null A_j, then $f \preceq g$, and moreover, if any of the $\preceq_{A_j}$ orderings are strict, then so is the $\preceq$ one.

Proposition IX.A.1. *Standard Dynamic Consistency on Ω is equivalent to DC.*

Note that this result can be thought of giving another argument that DC is the appropriate equivalent to SDC for the case of epistemic contingencies.

IX.B Proofs

Throughout the Appendix, B will denote the space of real-valued functions on S, and $ba(S)$ will denote the set of real-valued set functions on S, both under the Euclidean topology. B is equipped with the standard order: $a \leqslant b$ iff $a(s) \leqslant b(s)$ for all $s \in S$. For $x \in \Re$, x^* is the constant function taking value x; we use $\Re^*$ to denote the set of constant functions. Addition with acts and positive scalar multiplication is extended to sets as standard: for $A \subseteq B$, $a \in B$, $\alpha > 0$, $\alpha A = \{\alpha b \mid b \in A\}$ and $A + a = \{b + a \mid b \in A\}$. Finally, for any subset $A \subseteq B$, we shall use $co(A)$ to denote the convex closure of A.

IX.B.1 Proofs of Propositions

For the proofs of Propositions IX.1–IX.4, let $J : \mathcal{A} \to \Re$ be the maxmin EU functional represented by $\mathcal{C}$ (ie. $J(f) = \min_{p\in\mathcal{C}} \sum_{s\in S} u(f(s))p(s)$ for all $f \in \mathcal{A}$), and likewise for J_i and $\mathcal{C}_i$. $Ra(J_I)$ is the range of the vector $(J_i)_{i\in I}$ over $\mathcal{A}$, and $Ra(J_i)$ the range of the function J_i for each i. Note that, since the utility functions are the same for all $\preceq_i$, $Ra(J_i) = Ra(J_j)$ for all i, j; call this set R. $cone(Ra(J_I))$ is the cone spanned by $Ra(J_I)$. We use $\mathbf{e}$ to denote the unit vector in $\Re^I$.

Proofs of Proposition IX.1 and Proposition IX.2. Proposition IX.1 and the 'if' direction Proposition IX.2 are straightforward, and their proofs are omitted.

The proof of the 'only if' direction of Proposition IX.2 draws on the developments in Crès et al. (2011). By their Lemmas 1–4, there exists a constant additive, positively homogeneous, monotonic real-valued function ϕ on $cone(Ra(J_I))$ such that $J(f) = \phi((J_i(f))_{i\in I})$ for all $f \in \mathcal{A}$. Cerreia-Vioglio et al. (2013, Theorem 1) show that the real-valued function $\bar{\phi}$ on $\Re^I$, defined by $\hat{\phi}(\mathbf{y}) = \sup\{\phi(\mathbf{x}) + b \mid \mathbf{x} \in Ra(J_I),\ b \in \Re,\ \mathbf{x} + b\mathbf{e} \leqslant \mathbf{y}\}$ for all $\mathbf{y} \in \Re^I$, extends ϕ and is constant additive and monotonic. It is clear from the definition and the positive homogeneity of ϕ that $\hat{\phi}$ is positively homogeneous. Finally, strong monotonicity on $Ra(J_I)$ is a direct consequence of DC.

□

Proof of Proposition IX.3. The right to left direction is essentially Proposition IX.1. As for the left to right direction, by Proposition IX.2, ϕ is positively homogeneous, constant additive and monotonic. Applying (Hill, 2012, Lemmas 4 and 7), it follows that ϕ is concave, so it satisfies all the properties required for the Gilboa and Schmeidler (1989) agrument. Applying this reasoning (see also Crès et al. (2011)), we obtain a unique closed convex Λ representing ϕ as described in the Proposition. Finally, since, by (Hill, 2012, Lemma 4), $Ra(J_i)$ is fully dimensional, it follows from the strong monotonicity of ϕ on $Ra(J_I)$ and its positive homogeneity that ϕ is strongly monotonic, so every $p \in \mathcal{D}$ is such that $p(i) > 0$ for all $i \in I$.

□

Proof of Proposition IX.4. As shown by Crès et al. (2011, Lemmas 1–4), DC implies that $J(f) = \phi((J_i(f))_{i \in I})$ where ϕ is a constant additive, positively homogeneous, monotonic real-valued function on $Ra(J_I)$. Consider any vector $\mathbf{x} \in R^I$. For each $i \in I$, since $x_i \in R = Ra(J_i)$, there exists a constant act $c_i^{\mathbf{x}} \in \Delta(X)$ with $J_i(c_i^{\mathbf{x}}) = x_i$. Hence, defining $g^{\mathbf{x}} \in \mathcal{A}$ by $g(s) = c_i^{\mathbf{x}}(s)$ whenever $s \in A_i$, it follows from the fact that $\{\preceq_i\}_{i \in I}$ is $\mathcal{P}$-ontic that $((J_i(g^{\mathbf{x}}))_{i \in I}) = ((J_i(c_i^{\mathbf{x}}))_{i \in I}) = \mathbf{x}$. It follows that $R^I \subseteq Ra(J_I)$; by the definition of $Ra(J_I)$, this inclusion is in fact an equality. Now consider any pair of vectors $\mathbf{x}, \mathbf{y} \in R^I$. By definition, $J_i(\alpha g^{\mathbf{x}} + (1-\alpha) g^{\mathbf{y}}) = J_i(\alpha c_i^{\mathbf{x}} + (1-\alpha) c_i^{\mathbf{y}}) = \alpha J_i(c_i^{\mathbf{x}}) + (1-\alpha) J_i(c_i^{\mathbf{y}}) = \alpha J_i(g^{\mathbf{x}}) + (1-\alpha) J_i(g^{\mathbf{y}}) = \alpha x_i + (1-\alpha) y_i$, where the middle equality holds by the definition of the maxmin EU functional. Hence $\phi(\alpha \mathbf{x} + (1-\alpha)\mathbf{y}) = \phi((J_i(\alpha g^{\mathbf{x}} + (1-\alpha) g^{\mathbf{y}}))_{i \in I}) = J(\alpha g^{\mathbf{x}} + (1-\alpha) g^{\mathbf{y}}) \geqslant \alpha J(g^{\mathbf{x}}) + (1-\alpha) J(g^{\mathbf{y}}) = \alpha \phi((J_i(g^{\mathbf{x}}))_{i \in I}) + (1-\alpha)\phi((J_i(g^{\mathbf{y}})_{i \in I}) = \alpha\phi(\mathbf{x}) + (1-\alpha)\phi(\mathbf{y})$, where the inequality in the middle holds because of the concavity of the maxmin EU functional J. Hence ϕ is concave.

Using standard arguments, ϕ can be extended to a monotonic, positively homogeneous, constant additive, concave function on $\Re^I$. Application of the argument in Gilboa and Schmeidler (1989) (see also (Crès et al., 2011, Lemma 8)) implies that there exists a closed convex $\mathcal{D} \subseteq \Delta(I)$ such that $\phi(\mathbf{x}) = \min_{p \in \mathcal{D}} \sum_{i \in I} p_i x_i$. The strict positivity of the elements in $\mathcal{D}$ follows directly from DC and the fact that $Ra(J_I) = R^n$. The form of the set of priors $\mathcal{C}$ representing J follows from Crès et al. (2011, Proposition 1).

□

Proof of Proposition IX.5. By Proposition IX.2, DC implies that the aggregator ϕ is monotonic. Consider any $A \in \wp(\mathcal{A})$, and let $g = \arg\max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s) = \arg\max_{f\in A}\phi\left(\left(\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(f(s))p(s)\right)_{i\in I}\right)$. For every $i \in I$, by definition, $\max_{f\in A}\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(f(s))p(s) \geqslant \min_{p\in\mathcal{C}_i}\sum_{s\in S}u(g(s))p(s)$; it follows from the monotonicity of ϕ that $\phi\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(f(s))p(s)\right)_{i\in I}\right) \geqslant \phi\left(\left(\min_{p\in\mathcal{C}_i}\sum_{s\in S}u(g(s))p(s)\right)_{i\in I}\right) = \max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s)$, as required.

□

Proof of Proposition IX.6. To show the right to left direction, suppose that $\mathcal{C}_{(\mathcal{I})} \subseteq \mathcal{C}$. Take any $A \in \wp(\mathcal{A})$ let $g = \arg\max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s)$. By the definition of the maxmin EU rule (representation (IX.1)), the containment of sets of priors implies that $\min_{p\in\mathcal{C}_{(\mathcal{I})}}\sum_{s\in S}u(g(s))p(s) \geqslant \min_{p\in\mathcal{C}}\sum_{s\in S}u(g(s))p(s) = \max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s)$; however, it follows from Proposition IX.5 that $\phi_{(\mathcal{I})}\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u(f(s))p(s)\right)_{k\in K}\right) \geqslant \min_{p\in\mathcal{C}_{(\mathcal{I})}}\sum_{s\in S}u(g(s))p(s)$. Combining these two inequalities, one obtains that the value of information $\mathcal{I}$ is always non-negative.

Now consider the other direction, and suppose that $\mathcal{C}_{(\mathcal{I})} \nsubseteq \mathcal{C}$; we shall show that there exists $A \in \wp(\mathcal{A})$ with $\phi_{(\mathcal{I})}\left(\left(\max_{f\in A}\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u(f(s))p(s)\right)_{k\in K}\right) < \max_{f\in A}\min_{p\in\mathcal{C}}\sum_{s\in S}u(f(s))p(s)$. Since $\mathcal{C}_{(\mathcal{I})} \nsubseteq \mathcal{C}$, there exists $p \in \Delta(\Sigma)$ with $p \in \mathcal{C}_{(\mathcal{I})}\backslash\mathcal{C}$. By a separation theorem (Aliprantis and Border, 2007, 5.80), there is a nonzero linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ such that $\phi(p) \leqslant \alpha < \phi(q)$ for all $q \in \mathcal{C}$. Since S is finite (so B is finite-dimensional), B is reflexive, and, by the standard isomorphism between $ba(S)$ and B^*, it follows that $ba(S)^*$ is isometrically isomorphic to B (Dunford and Schwartz, 1958, IV.3); hence there is a real-valued function $a \in B$ such that $\phi(q) = \sum_{s\in S}a(s)p(s)$ for any $q \in ba(S)$. Without loss of generality ϕ, a can be chosen so that a takes values in the range of u. Take $g \in \mathcal{A}$ such that $u \circ g = a$, and consider the menu $\{g\}$. $\phi_{(\mathcal{I})}\left(\left(\max_{f\in\{g\}}\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u(f(s))p(s)\right)_{k\in K}\right) = \phi_{(\mathcal{I})}\left(\left(\min_{p\in\mathcal{C}_{(\mathcal{I})k}}\sum_{s\in S}u(g(s))p(s)\right)_{k\in K}\right) = \min_{p\in\mathcal{C}_{(\mathcal{I})}}\sum_{s\in S}u(g(s))p(s)$. However, by the definition of g, $\min_{p\in\mathcal{C}_{(\mathcal{I})}}\sum_{s\in S}u(g(s))p(s) <$

$\min_{p \in \mathcal{C}} \sum_{s \in S} u(g(s))p(s) = \max_{f \in \{g\}} \min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s)$. Hence $\phi_{(\mathcal{I})}\left(\left(\max_{f \in \{g\}} \min_{p \in \mathcal{C}_{(\mathcal{I})k}} \sum_{s \in S} u(f(s))p(s)\right)_{k \in K}\right) < \max_{f \in \{g\}} \min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s)$, as required.

□

Proof of Proposition IX.A.1. By the definition of learnable events, for each A_j, there exists $\mathcal{C}_j$ such that $\preceq_{A_j}$ is represented by $\mathcal{C}_j$ according to (IX.1). By the definition of Ω-nullness, for any Ω-non-null A_j, $\mathcal{C}_j \in \{\mathcal{C}_i\}_{i \in I}$ (it is one of the epistemic contingencies). Moreover, for any $A_{j_1}, A_{j_2} \in \{A_j\}_{j \in J}$, $\mathcal{C}_{j_1} \neq \mathcal{C}_{j_2}$: if not, then for any state s such that $p(s) > 0$ for some $p \in \mathcal{C}_j = \mathcal{C}_{j_1} = \mathcal{C}_{j_2}$, $(s, \mathcal{C}_j) \in A_{j_1} \cap A_{j_2}$, contradicting the disjointness of the elements of the partition. So, for any learnable partition $\{A_j\}_{j \in J}$, there is an injective map from the non-null elements of $\{A_j\}_{j \in J}$ to the set of epistemic contingencies $\{\mathcal{C}_i\}_{i \in I}$. Moreover, this map is surjective: for any epistemic contingency $\mathcal{C}_i$, there is a Ω-non-null state $(s, \mathcal{C}_i)$, and hence there is a cell $A_{j'}$ of the partition such that the associated set of priors $\mathcal{C}_{j'} = \mathcal{C}_i$. So, for any learnable partition, there is a bijection between the cells of this partition and the set of epistemic contingencies; moreover, this bijection is such that, for each cell of the partition, the preferences conditional on this cell coincide with the epistemic contingency related to it by the bijection. The equivalence between the two conditions follows immediately.

□

IX.B.2 Proof of Theorem IX.1

We first consider the direction (i) to (ii). The proof proceeds as follows. Firstly, using monotonicity (IX-A7) we show that c generates a choice function on a subset of the set of real-valued functions on S, which can be extended to a choice function on the whole set satisfying Chernoff, Aizerman, upper hemicontinuity and choice-function versions of constant linearity, superadditivity, monotonicity and union C-independence (Lemmas IX.B.1 and IX.B.2). The most important step in the proof is Proposition IX.B.1, which shows that, for any set $\hat{A}$ and any $\hat{a} \in c(\hat{A})$, there exists a closed convex set of finitely additive probability measures such that the maxmin expectation represents c in the sense that any maximal element over A' according to maxmin expectation with this set of priors is in $c(A')$, and such that $\hat{a}$ maximises the maxmin expctation on $\hat{A}$. Taking the union of such sets of probability measures over all pairs a, A with $a \in c(A)$ yields the required set $\mathcal{P}$.

IX.B.2.1 Preliminary Lemmas

By IX-A1 and IX-A4 and standard choice theory results (for example Sen (1971)), the restriction of c to sets containing only constant acts is represented in the standard way by a reflexive transitive complete order. By IX-A3, IX-A5 and IX-A9 and the Herstein-Milnor theorem, this order (and hence the restriction of the choice function) is represented by a non-degenerate affine utility function u; by IX-A9, u is continuous. Let $K = u(\Delta(X))$ and $B(K)$ be the set of functions in B taking values in K. Without loss of generality, it can be assumed that $0 \in K$, and it is not on its boundary.

Lemma IX.B.1. *For $A, B \subseteq \mathcal{A}$, suppose that there is a bijection $\sigma : A \to B$ such that, for all $f \in A$, $c(\{f(s), \sigma(f)(s)\}) = \{f(s), \sigma(f)(s)\}$ for all $s \in S$. Then $c(B) = \sigma(c(A))$.*

Proof. Let A and B satisfy the properties specified, and let $f \in c(A)$. By applying IX-A7 on the sets A and $A \cup B$, we have that $c(A) \subseteq c(A \cup B)$. We distinguish two cases. If there exists $d \in \Delta(X)$ with $d \notin c(\{f(s), d\})$ for all $s \in S$, then consider $c((A \cup B)_\alpha d \cup \{\sigma(f)\})$ for $\alpha \in (0, 1)$. By IX-A7, $f_\alpha d \notin c((A \cup B)_\alpha d \cup \{\sigma(f)\})$, whence $c(A)_\alpha d \nsubseteq c((A \cup B)_\alpha d \cup \{\sigma(f)\})$; it thus follows from IX-A2 and IX-A5 that $\sigma(f) \in c((A \cup B)_\alpha d \cup \{\sigma(f)\})$. By IX-A9 and the fact that, as $\alpha \to 1$, $(A \cup B)_\alpha d \cup \{\sigma(f)\} \to A \cup B$, $\sigma(f) \in c(A \cup B)$. By IX-A1, $\sigma(f) \in c(B)$ as required. If there is no $d \in \Delta(X)$ with $d \notin c(\{f(s), d\})$ for all $s \in S$, then take any $e \in \Delta(X)$ with $u(e)$ non-minimal in K and $\beta \in (0, 1)$. There exists $d \in \Delta(X)$ with $d \notin c(\{f_\beta e(s), d\})$ for all $s \in S$; applying the previous argument to $((A \cup B)_\beta e)_\alpha d \cup \{(\sigma(f))_\beta e\}$ yields the conclusion that $(\sigma(f))_\beta e \in c(B_\beta e)$. It follows by IX-A5 that $\sigma(f) \in c(B)$ as required. By a similar argument on $g \in c(B)$, the result is obtained.

□

There is thus a many-to-one mapping between acts in $\mathcal{A}$ and elements of $B(K)$, given by $a = u \circ f$, for $f \in \mathcal{A}$. Define the choice function $c_{B(K)}$, on $B(K)$, as follows: for $A \in \wp(B(K))$, and $A' \in \wp(\mathcal{A})$ such that $A = u \circ A'$, $c_{B(K)}(A) = u \circ c(A')$ By Lemma IX.B.1, $c_{B(K)}$ is well-defined. Let $\wp_{bbd}(B)$ be the set of closed bounded subsets of B.

Lemma IX.B.2. *There exists a choice function $c : \wp_{bdd}(B) \to \wp_{bdd}(B)$ on B such that:*

i. for all $A \in \wp(B(K))$, $c(A) = c_{B(K)}(A)$

ii. c satisfies Chernoff and Aizerman (that is, IX-A1 and IX-A2)[27]

27. Henceforth, we shall refer to these properties by these names.

iii. *c is constant linear: for all $A \in \wp_{bdd}(B)$, $\alpha > 0$ and $x \in \Re$, $c(\alpha A + x^*) = \alpha c(A) + x^*$*

iv. *c is constant independent: for all $A \in \wp_{bdd}(B)$, $\alpha \in (0,1)$ and $x \in \Re$ with $x^* \in c(A)$, $c(A) \subseteq c(A \cup (\alpha A + (1-\alpha)x^*))$*

v. *c is monotonic: for all $A, B \in \wp_{bdd}(B)$ with $A \subseteq B$, if for each $b \in B$, there exists $a \in A$ with $a \geqslant b$, then $c(A) \subseteq c(B)$, and, for any $b \in B$, if there exists $a \in B$ with $a > b$, then $b \notin c(B)$.*

vi. *c is superadditive: for all $A \in \wp_{bdd}(B)$, $A' \subseteq c(A)$, $a' \in c(A \cup \{a'\})$ for all $a' \in co(A')$ whenever there exist $x \in \Re$ and $b \in A'$ such that: i. $\alpha a + (1-\alpha)x^* \in A'$ for all $a \in A'$, $\alpha \in (0,1)$; and ii. $\alpha b + (1-\alpha)x^* \notin c(A' \cup \{a + y^*\})$ for all $a \in A'$, $y > 0$ and $\alpha \in (0,1)$.*

vii. *c is upper hemicontinuous: for all sequences $(A_n)_{n \in N}$, $A_n \in \wp_{bdd}(B)$ and $A \in \wp_{bdd}(B)$ with $A_n \to A$ and for all sequences $(a_n)_{n \in N}$, $a_n \in B$, with $a_n \in A_n$ for all $n \in N$, if $a_n \to a$, then $a \in c(A)$.*

Proof. Define c on $B(K)$ by clause i. Note that, applying IX-A5 on the inverse image of A and the inverse image of 0^*, we have that, for any $A \in \wp(B(K))$ and $\alpha \in (0,1)$, $c(\alpha A) = \alpha c(A)$. It follows that, for any $A \in \wp(B(K))$ with $\alpha A \in \wp(B(K))$ where $\alpha > 1$, $c(\alpha A) = \alpha c(A)$. c can thus be coherently extended to $\wp_{bdd}(B)$ as follows: for $A \in \wp_{bdd}(B)$, $c(A) = \frac{1}{\alpha} c(\alpha A)$, where $\alpha > 0$ is such that $\alpha A \in \wp(B(K))$. Note that c is positively homogeneous and satisfies the choice properties (point ii) by IX-A1 and IX-A2. Moreover, it is constant additive: applying IX-A5 to the inverse image of $2A$ and $2x^*$ (or appropriate products with a sufficiently small α), we have that $c(A + x^*) = c(\frac{1}{2}(2A) + \frac{1}{2}(2x^*)) = \frac{1}{2}c(2A) + \frac{1}{2}(2x^*) = c(A) + x^*$, as required. So c is constant linear. The remaining properties are a direct consequence of axioms IX-A6–IX-A9 (multiplying by a sufficiently small α where appropriate).

□

IX.B.2.2 Proposition IX.B.1

As stated above, the following proposition is the central part of the proof.

Proposition IX.B.1. *Let $c : \wp_{bdd}(B) \to \wp_{bdd}(B)$ be a choice function satisfying the properties in Lemma IX.B.2, and suppose that, for some $\hat{a} \in \hat{A} \in \wp_{bdd}(B)$, $\hat{a} \in c(\hat{A})$. Then there exists a closed convex set $\mathcal{C}_{\hat{a},\hat{A}}$ of finitely additive probability measures such that:*

(1) For all $A \in \wp_{bdd}(B)$,

(IX.9) $$b \in \arg\max_{b' \in A} \min_{p \in \mathcal{C}_{\hat{a},\hat{A}}} \sum_{s \in S} b'(s)p(s) \Rightarrow b \in c(A)$$

(2) $\hat{a} \in \arg\max_{b' \in \hat{A}} \min_{p \in \mathcal{C}_{\hat{a},\hat{A}}} \sum_{s \in S} b'(s)p(s)$

Proof. We consider the case where $\hat{a}$ is not a constant act; the case where is it a constant act is treated similarly, by using the construction below with a non-constant act a such that $\{\alpha a + (1-\alpha)\hat{a} |\ \alpha \in [0,1]\} = c(\{\alpha a + (1-\alpha)\hat{a} |\ \alpha \in [0,1]\})$.

We begin with some notation. First, recall that B is (isomorphic to) Euclidean space, so for ease we may use Euclidean notation and intuition at points. In particular, let $\|\cdot\|$ be the Euclidean norm. Let $U_{\hat{a}} = \{b \in B \mid \inf_{x \in \Re} \|b - x^*\| \leqslant \inf_{x \in \Re} \|\hat{a} - x^*\|\}$: this is the smallest 'tube' around the ray generated by the unit vector containing $\hat{a}$. By the assumption that $\hat{a}$ is not constant, it is not a ray. Moreover, for any $a \in B$, let $\overline{a} = \{\alpha a + x^* |\ \alpha \geqslant 0,\ x \in \Re\}$; this is the positive half-plane generated by a and the unit vector, with as boundary the ray generated by the unit vector. It is straightforward to check that the $\{\overline{a} \backslash \Re^* |\ a \in B \backslash \Re^*\}$ form a partition of $B \backslash \Re^*$; call the set of equivalence classes of the partition $\mathcal{Q}$.

For $a \in B$, $x \in \Re$, we define $[a, x^*] = \{\alpha a + (1-\alpha)x^* |\ \alpha \in [0,1]\}$, and $\Xi = \{[a, x^*] |\ a \in B,\ x \in \Re,\ c([a, x^*]) = [a, x^*]\}$. Moreover, for any $A \in \wp_{bdd}(B)$ and $x \in \Re$, $X_x^A = \bigcup_{a \in A} [a, x^*]$. Note that, since A is closed and bounded, so is X_x^A.

Now consider the set $\mathcal{Z}$ of pairs (O, I) where:

1. O is a non-empty convex closed subset of B such that:
   a. for all $a \in B$, if $a \in O$, then $\overline{a} \subseteq O$
   b. $\hat{a} \in O$
2. $I : O \to \Re$ is a functional with the following properties:
   a. I is monotonic: for all $a, b \in B$, if $a \geqslant b$, $I(a) \geqslant I(b)$
   b. I is constant linear (ie. constant additive and positively homogeneous): for all $a \in O$, $x \in \Re$, $\alpha > 0$, $I(\alpha a + x^*) = \alpha I(a) + x$
   c. I is superadditive: for any $\alpha \in [0,1]$, $I(\alpha a + (1-\alpha)b) \geqslant \alpha I(a) + (1-\alpha)I(b)$
   d. I represents c on O: for all $A \in \wp_{bdd}(O)$, $\arg\max_A I \subseteq c(A)$

e. $\hat{a} \in \arg\max_{\hat{A}\cap O} I$

f. for any $A' \in \wp_{bdd}(O)$ with $A' \subseteq \{a \in O | I(a) = I(\hat{a})\}$, $A' \subseteq c(A' \cup X^{\hat{A}}_{I(\hat{a})})$

g. for every $A \in \wp_{bdd}(O)$ with $(\{a \in O | I(a) = I(\hat{a})\} \cap U_{\hat{a}}) \cup X^{\hat{A}}_{I(\hat{a})} \subseteq A$, and for every $a' \in \{a \in O | I(a) = I(\hat{a})\} \cap U_{\hat{a}}$ and every $z > 0$, $\alpha\hat{a} + (1-\alpha)(I(\hat{a}))^* \notin c(A \cup \{a' + z^*\})$ for every $\alpha \in (0,1)$.

Note that, by property 2f, Chernoff, and the definition of $U_{\hat{a}}$, $[\hat{a}, I(\hat{a})^*] \subseteq \{a \in O | I(a) = I(\hat{a})\} \cap U_{\hat{a}}$.

$\mathcal{Z}$ is equipped with the order $\leqslant$, defined as follows: $(O_1, I_1) \leqslant (O_2, I_2)$ iff:

— $O_1 \subseteq O_2$

— $I_1 = I_2|_{O_1}$

This is evidently a partial order. In a pair of auxiliary lemmas (proved in Section IX.B.3), we show that any element of this order not containing some $b \in B$ can be extended to another element in $\mathcal{Z}$ containing it (Lemma IX.B.4) and that $\mathcal{Z}$ is non-empty (Lemma IX.B.12). Since, for each $(O, I) \in \mathcal{Z}$, O is the convex hull of a set of half-planes, and B is finite dimensional, all chains of elements of $\mathcal{Z}$ must be finite. Hence every chain has an upper bound, which is in fact its top element. Take any maximal chain, in the sense that there is no $(O'', I'') \in \mathcal{Z}$ which is $(O'', I'') > (O', I')$ for all (O', I') in the chain. For the top element of this chain, $O = B$; if not, then there exists $b \notin O$, whence, by Lemma IX.B.4, there is a $(O', I') \in \mathcal{Z}$ with $(O', I') > (O, I)$. Hence there is an element $(B, I) \in \mathcal{Z}$, with I being a monotonic, constant linear, superadditive functional on B representing c (in the sense of property 2d) and satisfying property 2e.

By the result in Gilboa and Schmeidler (1989) (in particular Lemma 3.5 onwards), there exists a unique closed convex set of probability measures $\mathcal{C}_{\hat{a},\hat{A}}$ such that $I(a) = \min_{c\in\mathcal{C}_{\hat{a},\hat{A}}} \sum_{s\in S} a(s)p(s)$. By the fact that I represents c, it follows that (IX.9) holds for $\mathcal{C}_{\hat{a},\hat{A}}$. Moreover, by property 2e, $\hat{a} \in \arg\max_{\hat{A}} I$, and so $\hat{a} \in \arg\max_{b\in\hat{A}} \min_{p\in\mathcal{C}_{\hat{a},\hat{A}}} \sum_{s\in S} b(s)p(s)$, as required.

□

IX.B.2.3 Conclusion of the proof of Theorem IX.1

We conclude the direction (i) to (ii). Let $\mathcal{P} = \{\mathcal{C}_{\hat{a},\hat{A}} \mid \hat{A} \in \wp_{bdd}(B),\ \hat{a} \in c(\hat{A})\}$, where the $\mathcal{C}_{\hat{a},\hat{A}}$ are as in Proposition IX.B.1. By the definition of the choice function c on B, and clause 1 of Proposition IX.B.1, for any $\mathcal{C} \in \mathcal{P}$, if $f \in \arg\max_{g\in A} \min_{p\in\mathcal{C}} \sum_{s\in S} u(g(s))p(s)$, then $f \in c(A)$. Moreover, for any $f \in \mathcal{A}$ and $A \in \wp(\mathcal{A})$, if $f \in c(A)$, then, by clause 2 of Proposition IX.B.1, $f \in \arg\max_{g \in A} \min_{p \in \mathcal{C}_{u \circ f, u \circ A}} \sum_{s \in S} u(g(s))p(s)$. Hence $\mathcal{P}$ and u represent c according to (IX.5).

As concerns the necessity of the axioms (the (ii) to (i) direction), all axioms are evident or have been shown to be necessary elsewhere in the literature (see for example Moulin (1985) for the necessity of IX-A1 and IX-A2), except IX-A8. To establish the necessity of this axiom, we first claim that the conditions of the axiom imply that there is a set $\mathcal{C} \in \mathcal{P}$ such that $B \subseteq \arg\max_{h \in A} \min_{p \in \mathcal{C}} \sum_{s \in S} u(h(s))p(s)$. Suppose that this is not the case and consider the act $g \in B$ mentioned in the axiom. Take any set $\mathcal{C}_1 \in \mathcal{P}$ such that $g_\alpha d \in \arg\max_{h \in A} \min_{p \in \mathcal{C}_1} \sum_{s \in S} u(h(s))p(s)$ for some $\alpha \in (0, 1)$. By the reductio assumption, there must be an act $f \in B$ such that $f \notin \arg\max_{h \in A} \min_{p \in \mathcal{C}_1} \sum_{s \in S} u(h(s))p(s)$; take such an f. By the facts just established, and the fact that $d \in A$, we have that $\min_{p \in \mathcal{C}_1} \sum_{s \in S} u(g_\alpha d(s))p(s) \geqslant u(d)$ and $\min_{p \in \mathcal{C}_1} \sum_{s \in S} u(g_\alpha d(s))p(s) > \min_{p \in \mathcal{C}_1} \sum_{s \in S} u(f(s))p(s)$. It thus follows that there exists $e \succ d$ and $\beta \in (0, 1)$ such that $\min_{p \in \mathcal{C}_1} \sum_{s \in S} u(g_\alpha d(s))p(s) \geqslant \min_{p \in \mathcal{C}_1} \sum_{s \in S} u(f_\beta e(s))p(s)$, and hence, by the representation (IX.5), that $g_\alpha d \in c(A \cup \{f_\beta e\})$, contradicting IX-A8. Hence there is a set $\mathcal{C} \in \mathcal{P}$ such that $B \subseteq \arg\max_{h \in A} \min_{p \in \mathcal{C}} \sum_{s \in S} u(h(s))p(s)$, as required. It follows by the concavity of the maxmin expected utility representation that, for any $f \in conv(B)$, $\min_{p \in \mathcal{C}} \sum_{s \in S} u(f(s))p(s) \geqslant \max_{h \in A} \min_{p \in \mathcal{C}} \sum_{s \in S} u(h(s))p(s)$, and so by the representation (IX.5), it follows that $f \in c(A \cup \{f\})$ for all such f, as required.

The uniqueness of u follows from the standard von Neuman-Morgenstern result. We now show the existence of a unique minimal $\mathcal{P}$ representing c. Let $\{\mathcal{P}_m\}_{m \in M}$ be the sets of sets of priors representing c according to (IX.5) and let $\mathcal{P} = \bigcup_{m \in M} \mathcal{P}_m$. Let I index $\mathcal{P}$. Note that $\mathcal{P}$ also represents c according to (IX.5). Pick any $d \in \Delta(X)$ and for each $\mathcal{C}_i \in \mathcal{P}$, let $A_i = \{f \in \mathcal{A} \mid \min_{p \in \mathcal{C}_i} \sum_{s \in S} u(f(s))p(s) = u(d)\}$. Since $\mathcal{P}$ represents c according to (IX.5), $c(A_i) = A_i$ for all A_i. Define the following order on $\{A_i \mid i \in I\}$: for all $i, j \in I$, $A_i \sqsupseteq A_j$ iff there exists $f \in A_i$, $e \in \Delta(X)$ with $e \prec d$ and $\alpha \in (0, 1)$ such that $f_\alpha e \in A_j$. Note that, by the definition of A_i, and the fact that, for all $i, j \in I$, $A_i \neq A_j$, $\sqsubseteq$ is complete, and the strict relation is transitive. We say that a set A_i is *essential* if there exists no set $J \subset I \backslash \{i\}$ such that $A_i \not\sqsupseteq A_j$ for all $j \in J$ and $\bigcup_{j \in J} A_j \supseteq A_i$.

To get a preliminary understanding of these notions, note that, for any $i \in I$, if $f \in A_i$ then $f_\alpha d \in A_i$ for all $\alpha \in [0, 1]$. Moreover, it follows from the representation that, if

$A_i \sqsupseteq A_j$, then there exists $f \in A_i$ such that $f >_j d$. Hence, by the properties of the maxmin EU representation, for all $\alpha \in [0,1)$, $f_\alpha d \notin \{h \in A_i \mid h \geq_j g,\ \forall g \in A_i\}$. Moreover, once again by the properties of the EU maxmin representation, we have that, for any act f' with $f'_\beta d \in A_i \cap A_j$ for all $\beta \in [0,1]$, $f'_\beta d \notin \{h \in A_i \mid h \geq_j g,\ \forall g \in A_i\}$ for all $\beta \in [0,1)$. It follows that not only $\{h \in A_i \mid h \geq_j g,\ \forall g \in A_i\} \neq A_i$ for any A_j such that $A_i \sqsupseteq A_j$, but also that $\bigcup_{j \text{ s.t. } A_i \sqsupseteq A_j} \{h \in A_i \mid h \geq_j g,\ \forall g \in A_i\} \neq A_i$. This motivates the definition of essential sets: they are those sets A_i for which the fact that $A_i = c(A_i)$ cannot be 'attained' using a set of sets A_j. Uniqueness follows from Lemma IX.B.3 below, which implies that there is a unique minimal subset of $\mathcal{P}$ representing c; it follows in particular that this is contained in every set of sets of priors $\mathcal{P}_m$ representing c.

Lemma IX.B.3. *$\{\mathcal{C}_i \mid i \in I,\ A_i$ essential$\}$ represents c according to (IX.5), and for any set $\{\mathcal{C}_k \mid k \in K\}$ with $K \subseteq I$ that represents c, $K \supseteq \{i \in I \mid A_i$ essential$\}$.*

Proof. For ease of presentation, we reason on the preference relations $\preceq_i$ generated by the sets of priors $\mathcal{C}_i$; recall that the corresponding representation in terms of preference orderings is (IX.6), and that we use $\mathcal{R}$ to denote the set of preference relations generated by $\mathcal{P}$. We first show that $\{\preceq_i \mid i \in I,\ A_i$ essential$\}$ represents c. Since $\{\preceq_i \mid i \in I\} = \mathcal{R}$ represents c, for each $\preceq_i \in \{\preceq_i \mid i \in I,\ A_i$ essential$\}$ and each set $A \in \wp(\mathcal{A})$, $\{f \in A \mid f \succeq_i g,\ \forall g \in A\} \subseteq c(A)$. So for every $A \in \wp(\mathcal{A})$, $\bigcup_{i \text{ s.t.} A_i \text{ essential}} \{f \in A \mid f \succeq_i g,\ \forall g \in A\} \subseteq c(A)$. It remains to show the inverse inclusion. For reductio, suppose that it does not hold, that is, that there exists $A \in \wp(\mathcal{A})$ and $f' \in A$ such that $f' \notin \bigcup_{i \text{ s.t.} A_i \text{ essential}} \{f \in A \mid f \succeq_i g,\ \forall g \in A\}$ but $f' \in c(A)$. By the properties of the maxmin EU functional, the fact that $\mathcal{R}$ represents c, and the definition of A_i, it follows that there exists $f \in \bigcup_{i \in I} A_i \backslash \bigcup_{A_i \text{ essential}} A_i$. Take any A_i with $f \in A_i$ that is $\sqsupseteq$-maximal – ie. for every A_k such that $f \in A_k$, $A_i \sqsupseteq A_k$. Since $\sqsupseteq$ is complete and the strict relation generated by it is transitive, such a maximal element exists. For such A_i, every set $J \subset I \backslash \{i\}$ such that $\bigcup_{j \in J} A_j \supseteq A_i$ must contain some k such that $f \in A_k$; however, by the definition of A_i, it follows that $A_i \sqsupseteq A_k$. Hence A_i is essential, contrary to the definition of f. So, by reductio, there exists no such f, and $\bigcup_{i \text{ s.t.} A_i \text{ essential}} \{f \in A \mid f \succeq_i g,\ \forall g \in A\} = c(A)$; so $\{\preceq_i \mid i \in I,\ A_i$ essential$\}$ represents c, as required.

We now show that any subset of $\mathcal{R}$ representing c must contain all elements yielding essential A_i. For reductio, suppose that there exists $K \subseteq I$ with $\{\preceq_k \mid k \in K \subseteq I\}$ representing c and $K \not\supseteq \{i \in I \mid A_i$ essential$\}$. Take any $i \in \{i \in I \mid A_i$ essential$\} \backslash K$, and consider A_i. Since $c(A_i) = A_i$ by definition, there must exist $J \subseteq K$ such that $\bigcup_{j \in J} A_j \supseteq A_i$ and $A_i \not\supseteq A_j$ for all $j \in J$. However, the existence of such a set contradicts the assumption that A_i is essential. Hence there exists no such K representing c, as required. Hence the claim is established.

□

IX.B.3 Auxiliary Lemmas for the Proof of Proposition IX.B.1

Lemma IX.B.4. *Let $(O, I) \in \mathcal{Z}$ be such that, for $b \in B$, $b \notin O$. Then there exists $(O', I') \geqslant (O, I)$ with $b \in O'$.*

Proof of Lemma IX.B.4. Since $b \notin O$, $\bar{b} \cap O = \Re^*$. Let $O' = co(O \cup \bar{b})$. We now extend I to I' on O'. Throughout the proof of this Lemma (and in particular Lemmas IX.B.5–IX.B.11), I will remain fixed; as a point of notation, for any $a \in O$, we let $x_a \in \Re$ be such that $I(a) = x_a$.

Let $[x_{\hat{a}}]^O = \{a \in O | I(a) = I(\hat{a})\} \cap U_{\hat{a}}$. Note that, by the monotonicity and constant linearity of I and the fact that $U_{\hat{a}}$ is closed, $[x_{\hat{a}}]^O$ is closed and bounded. Note also that, by the constant linearity of I, for any $a \in [x_{\hat{a}}]^O$ and $\beta \in [0, 1]$, $\beta a + (1 - \beta)x^*_{\hat{a}} \in [x_{\hat{a}}]^O$; hence, in particular, $\beta[x_{\hat{a}}]^O + (1 - \beta)x^*_{\hat{a}} \subseteq [x_{\hat{a}}]^O$.

Let $F_b = \overline{(O' \backslash O)} \cap U_{\hat{a}*} \cap \{a \in B \mid \sum_{s \in S} a(s) = x_{\hat{a}}\}$. (Geometrically, this is the intersection between the closure of $O' \backslash O$, $U_{\hat{a}*}$ and the hyperplane normal to the unit vector going through $x^*_{\hat{a}}$.) F_b is evidently closed and bounded and hence compact. Finally, for a compact subspace $F \subseteq B$ and a continuous bounded function $\sigma : F \to \Re$, we let $Y_\sigma = \{a + \sigma(a)^* \mid a \in F\}$. Note that Y_σ is closed and bounded, because it is the image of a continuous map from a compact space to a Hausdorff one. We use $\leqslant$ to denote the standard dominance order on functions σ ($\sigma \leqslant \sigma'$ iff $\sigma(a) \leqslant \sigma'(a)$ for all $a \in F$).

Now consider the following set:

$$C_b = \left\{ \sigma : F_b \to \Re \;\middle|\; \begin{array}{l} [x_{\hat{a}}]^O \cup Y_\sigma \subseteq c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}}\right) \ \& \\ \forall d \in Y_\sigma,\ [d, x^*_{\hat{a}}] \subseteq Y_\sigma \ \& \\ \forall a' \in [x_{\hat{a}}]^O \cup Y_\sigma,\ \forall z > 0, \forall \alpha \in (0, 1), \\ \alpha\hat{a} + (1 - \alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{a' + z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right) \end{array} \right\} \tag{IX.10}$$

Note that, since $[x_{\hat{a}}]^O$, Y_σ and $X^{\hat{A}}_{x_{\hat{a}}}$ are closed and bounded, so is $[x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}}$.

Lemma IX.B.5. $C_b \neq \varnothing$.

Proof. Consider $\{\sigma : F_b \to \Re | \ [x_{\hat{a}}]^O \subseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})\}$. This set is non-empty by property 2f of I and the monotonicity of c (property v in Lemma IX.B.2). Moreover, by the monotonicity of c and Chernoff, for any $\sigma', \sigma'' : F_b \to \Re$ with $\sigma' \geqslant \sigma''$ if σ' is in this set, then so is σ''. It follows by continuity (property vii in Lemma IX.B.2) that this set has a maximum element; let σ be any such element. By definition, we thus have that $[x_{\hat{a}}]^O \subseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$. Now consider any $c \in F_b$ and $\epsilon > 0$. By the maximality of σ, $[x_{\hat{a}}]^O \nsubseteq c([x_{\hat{a}}]^O \cup Y_{\sigma_c^{+\epsilon}} \cup X^{\hat{A}}_{x_{\hat{a}}})$, where $\sigma_c^{+\epsilon}(c) = \sigma(c) + \epsilon$ and $\sigma_c^{+\epsilon}(d) = \sigma(d)$ for $d \neq c$; by Chernoff, it follows that $[x_{\hat{a}}]^O \nsubseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{c + (\sigma(c)+\epsilon)^*\} \cup X^{\hat{A}}_{x_{\hat{a}}})$. It follows from Aizerman that $c + (\sigma(c)+\epsilon)^* \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{c + (\sigma(c)+\epsilon)^*\} \cup X^{\hat{A}}_{x_{\hat{a}}})$. Since this holds for any $\epsilon > 0$, it follows from continuity (property vii) that $c + \sigma(c)^* \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$. Since this holds for all $c \in F_b$, we have that $[x_{\hat{a}}]^O \cup Y_\sigma \subseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$. (Note that it follows in particular, using monotonicity, that for all $a \in [x_{\hat{a}}]^O \cap F_b$, $\sigma(a) = 0$.)

We now show that $[d, x^*_{\hat{a}}] \subseteq Y_\sigma$ for all $d \in Y_\sigma$. Take any $d \in Y_\sigma$ and any $\gamma \in (0,1)$, and consider $c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}} \cup (\gamma \left([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}}\right) + (1-\gamma)x^*_{\hat{a}}))$. Since, by their definition, $\gamma[x_{\hat{a}}]^O + (1-\gamma)x^*_{\hat{a}} \subseteq [x_{\hat{a}}]^O$ and $\gamma X^{\hat{A}}_{x_{\hat{a}}} + (1-\gamma)x^*_{\hat{a}} \subseteq X^{\hat{A}}_{x_{\hat{a}}}$, $c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}} \cup (\gamma \left([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}}\right) + (1-\gamma)x^*_{\hat{a}})) = c([x_{\hat{a}}]^O \cup Y_\sigma \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$. Since $x^*_{\hat{a}} \in [x_{\hat{a}}]^O$, it follows from constant independence (property iv) that $[x_{\hat{a}}]^O \cup Y_\sigma \subseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$. Moreover, it follows from the maximality of Y_σ that, for every $z > 0$, $[x_{\hat{a}}]^O \nsubseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d + z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}})$, so, by constant linearity $\gamma[x_{\hat{a}}]^O + (1-\gamma)x^*_{\hat{a}} \nsubseteq c(\gamma \left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d + z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right) + (1-\gamma)x^*_{\hat{a}})$, whence, by Chernoff $\gamma[x_{\hat{a}}]^O + (1-\gamma)x^*_{\hat{a}} \nsubseteq c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{\gamma(d + z^*) + (1-\gamma)x_{\hat{a}}\} \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$. By Aizerman it follows that $\gamma(d + z^*) + (1-\gamma)x^*_{\hat{a}} \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{\gamma(d + z^*) + (1-\gamma)x^*_{\hat{a}}\} \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$ for every $z > 0$, so by continuity $\gamma d + (1-\gamma)x^*_{\hat{a}} \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$. Since $d \in Y_\sigma$, it follows by the definition of Y_σ that there exists $x \in \Re$ such that $\gamma d + x^* \in Y_\sigma$. Hence, $\gamma d + x^*, \gamma d + (1-\gamma)x^*_{\hat{a}} \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup (\gamma Y_\sigma + (1-\gamma)x^*_{\hat{a}}) \cup X^{\hat{A}}_{x_{\hat{a}}})$; by monotonicity (property v), it follows that $x = (1-\gamma)x_{\hat{a}}$, and so $\gamma d + (1-\gamma)x^*_{\hat{a}} \in Y_\sigma$. Since this holds for every $\gamma \in (0,1)$, we have that $[d, x^*_{\hat{a}}] \subseteq Y_\sigma$, as required.

It remains to show that $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{a' + z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for all $a' \in [x_{\hat{a}}]^O \cup Y_\sigma$, $z > 0$ and $\alpha \in (0,1)$. First note that, by property 2g, $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{a' + z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for all $a' \in [x_{\hat{a}}]^O$, $z > 0$ and

$\alpha \in (0,1)$. We show that this is in fact the case for every $a' \in [x_{\hat{a}}]^O \cup Y_\sigma$. For reductio, take any $d \in Y_\sigma$ and $z > 0$ and $\alpha \in (0,1)$ and suppose that $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \in c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$. Consider any $a' \in [x_{\hat{a}}]^O$ and $\gamma \in (0,1)$; by property 2g, $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{\gamma(a'+y^*) + (1-\gamma)x^*_{\hat{a}}, d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for any $y > 0$. It follows from Aizerman that $\gamma(a'+y^*) + (1-\gamma)x^*_{\hat{a}} \in c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{\gamma(a'+y^*) + (1-\gamma)x^*_{\hat{a}}, d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$, and hence by continuity (property vii), $\gamma a' + (1-\gamma)x^*_{\hat{a}} \in c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$. Since this holds for every $\gamma \in (0,1)$ and $a' \in [x_{\hat{a}}]^O$, it follows that $[x_{\hat{a}}]^O \subseteq c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$, contradicting the maximality of σ. Hence $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for any $d \in Y_\sigma$, $z > 0$ and $\alpha \in (0,1)$, and so $\alpha\hat{a} + (1-\alpha)x^*_{\hat{a}} \notin c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{a'+z^*\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for all $a' \in [x_{\hat{a}}]^O \cup Y_\sigma$, $z > 0$ and $\alpha \in (0,1)$, as required.

□

Lemma IX.B.6. *For all $d \in F_b$, $\alpha > 0$ and $x \in \Re$ such that $\alpha d + x^* \in F_b$, and all $\sigma \in C_b$, $\sigma(\alpha d + x^*) = (1-\alpha)x^*_{\hat{a}} - x^* + \alpha\sigma(d)$.*

Proof. It suffices to show this for all $\alpha \in (0,1]$; the other cases follow immediately. Suppose that $d, \alpha d + x^* \in F_b$ for $\alpha \in (0,1]$, and $\sigma \in C_b$. Since $d + \sigma(d)^* \in Y_\sigma$, it follows from the definition of C_b that $\alpha(d+\sigma(d)^*) + (1-\alpha)x^*_{\hat{a}} \in Y_\sigma$. It thus follows by monotonicity (property v) and the definition C_b, $\sigma(\alpha d + x^*)$ is such that $\alpha d + x^* + \sigma(\alpha d + x^*)^* = \alpha(d + \sigma(d)^*) + (1-\alpha)x^*_{\hat{a}}$. So $\sigma(\alpha d + x^*) = (1-\alpha)x^*_{\hat{a}} - x^* + \alpha\sigma(d)$, as required.

□

Take any element $\sigma \in C_b$; Lemma IX.B.5 guarantees that such an element exists. Let $I'' : O \cup F_b \to \Re$ to be the functional extending I and such that $I''(d) = x_{\hat{a}} - \sigma(d)$ for $d \in F_b$. By Lemma IX.B.6, I'' is constant linear where defined; hence, by the definition of F_b and O', there exists a unique constant linear extension to O', which we call I'. By the definition of C_b, the constant linearity of c and Chernoff, I' satisfies properties 2f and 2g. We now establish several other properties of I'.

Lemma IX.B.7. *I' satisfies property 2e: for all $d \in O'$, if $d \in \hat{A}$, then $I'(d) \leqslant x_{\hat{a}}$.*

Proof. Since for $d \in O$ this follows from the properties of I, it suffices to consider $d \in O' \backslash O$. Suppose for reductio that $d \in \hat{A}$ and $I'(d) > x_{\hat{a}}$. By the definition of F_b, there exists

$\beta \in (0,1]$, $x \in \Re$ such that $\beta d + x^* \in F_b$; hence $\frac{1}{\beta}(x_{\hat{a}} - \sigma(\beta d + x^*) - x) = I'(d) > x_{\hat{a}}$. By the definition of C_b, we have that $\beta d + x^* + \sigma(\beta d + x^*)^* \in c([x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$; however, by the monotonicity of c (property v of Lemma IX.B.2), since $\beta d + (1-\beta)x^*_{\hat{a}} \in X^{\hat{A}}_{x_{\hat{a}}} \subseteq [x_{\hat{a}}]^O \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}}$ and $x + \sigma(\beta d + x^*) < (1-\beta)x_{\hat{a}}$, $\beta d + x^* + \sigma(\beta d + x^*)^* \notin c([x_{\hat{a}}]^O \cup [b + z^*, x^*_{\hat{a}}] \cup X^{\hat{A}}_{x_{\hat{a}}})$, which is a contradiction. So $I'(d) \leqslant x_{\hat{a}}$, as required.

□

Lemma IX.B.8. *I' is superadditive: for all $a, a' \in O'$ and $\alpha \in [0,1]$, then $I'(\alpha a + (1-\alpha)a') \geqslant \alpha I'(a) + (1-\alpha)I(a')$.*

Proof. We show the result for $a, a' \in U_{\hat{a}}$; it extends to other cases by the constant linearity of I'. Suppose for reductio that for some $a, a' \in O' \cap U_{\hat{a}}$, $\alpha \in [0,1]$, $I'(\alpha a + (1-\alpha)a') < \alpha I'(a) + (1-\alpha)I'(a')$. By the convexity of $U_{\hat{a}}$ and the definition of $[x_{\hat{a}}]^O$, F_b and I', $\alpha a + (1-\alpha)a' + (x_{\hat{a}} - I'(\alpha a + (1-\alpha)a'))^* \in [x_{\hat{a}}]^O \cup Y_\sigma$. Similarly, $a + (x_{\hat{a}} - I'(a))^*, a' + (x_{\hat{a}} - I'(a'))^* \in [x_{\hat{a}}]^O \cup Y_\sigma$. It follows from the definition of C_b and superadditivity (property vi of Lemma IX.B.2) that $d \in c\left([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d\} \cup X^{\hat{A}}_{x_{\hat{a}}}\right)$ for all $d \in co([x_{\hat{a}}]^O \cup Y_\sigma)$; hence $\alpha(a + (x_{\hat{a}} - I'(a))^*) + (1-\alpha)(a' + (x_{\hat{a}} - I'(a'))^*) \in c([x_{\hat{a}}]^O \cup \{\alpha(a + (x_{\hat{a}} - I'(a))^*) + (1-\alpha)(a' + (x_{\hat{a}} - I'(a'))^*)\} \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$. However, since $I'(\alpha a + (1-\alpha)a') < \alpha I'(a) + (1-\alpha)I(a')$, $\alpha a + (1-\alpha)a' + (x_{\hat{a}} - I'(\alpha a + (1-\alpha)a')))^* \in Y_\sigma$ strictly dominates $\alpha(a + (x_{\hat{a}} - I'(a))^*) + (1-\alpha)(a' + (x_{\hat{a}} - I'(a'))^*)$, and so it follows from that monotonicity of c (property v of Lemma IX.B.2) that $\alpha(a + (x_{\hat{a}} - I'(a))^*) + (1-\alpha)(a' + (x_{\hat{a}} - I'(a'))^*) \notin c([x_{\hat{a}}]^O \cup \{\alpha(a + (x_{\hat{a}} - I'(a)))^*) + (1-\alpha)(a' + (x_{\hat{a}} - I'(a')))^*)\} \cup Y_\sigma \cup X^{\hat{A}}_{x_{\hat{a}}})$, which is a contradiction. Hence $I'(\alpha a + (1-\alpha)a')) \geqslant \alpha I'(a) + (1-\alpha)I(a')$, as required.

□

Lemma IX.B.9. *I' is monotonic: for every $a, d \in O'$, if $d \leqslant a$, then $I'(d) \leqslant I'(a)$.*

Proof. We show the result for $a, d \in U_{\hat{a}}$; it extends to other cases by the constant linearity of I'. Consider $a, a' \in O' \cap U_{\hat{a}}$; by the definition of $[x_{\hat{a}}]^O$, F_b and I', $a + (x_{\hat{a}} - I'(a))^*, d + (x_{\hat{a}} - I'(d))^* \in [x_{\hat{a}}]^O \cup Y_\sigma$. If $d \leqslant a$, then, by the monotonicity of c, for each $y > 0$, since $d + (x_{\hat{a}} - I'(a) - y)^* < a + (x_{\hat{a}} - I'(a))^*$, $d + (x_{\hat{a}} - I'(a) - y)^* \notin c([x_{\hat{a}}]^O \cup Y_\sigma \cup \{d + (x_{\hat{a}} - I'(a) - y)^*\} \cup X^{\hat{A}}_{x_{\hat{a}}})$, and so $x_{\hat{a}} - I'(a) - y \neq x_{\hat{a}} - I'(d)$. Hence $I'(a) \geqslant I'(d)$, as required.

□

It remains to show that I' represents c on O' (property 2d). For this, we need the following preliminary lemma.

Lemma IX.B.10. *Let $A \in \wp_{bdd}(O')$, $A \subset P$ for some $P \in \mathcal{Q}$, and let $x \in \Re$ be such that $x \geqslant I'(a'')$ for all $a'' \in A$. Then there exists $a \in P$ such that $c(A \cup [a, x^*]) = [a, x^*]$.*

Proof. By the constant linearity of I', there exists $a' \in P$ with $I'(a') = x$. Moreover, for any $d \in A$, since $a', d \in P$, it follows from the definition of $\mathcal{Q}$ that $d = \alpha a' + z^*$ for some $\alpha \geqslant 0$ and $z \in \Re$. Hence there exists $\gamma_d \geqslant 1$, $\beta' \in [0, 1]$ and $y' \in \Re$ such that $d = \beta'(\gamma_d a' + (1 - \gamma_d)x^*) + (1 - \beta')y'^*$; take any such γ_d. Since A is closed and bounded, there exists $\gamma \geqslant \gamma_d$ for all $d \in A$; let $a = \gamma a' + (1 - \gamma)x^*$. By construction (and the constant linearity of I'), $I'(a) = x$, and, for each $d \in A$, there exists $\beta \in [0, 1]$ and $y \in \Re$ such that $d = \beta a + (1 - \beta)y^*$. It follows that, for any $d \in A$ with $I'(d) = x$, $d \in [a, x^*]$.

First we show that, for each $d \in A$ with $I'(d) < x$, $d \notin c(A \cup [a, x^*])$. Let $d \in A$ be such that $I'(d) < x$. By the previous observation, there exists $\beta \in [0, 1)$ and $y \in \Re$ such that $d = \beta a + (1 - \beta)y^*$. Since, by constant linearity of I', $I'(d) = \beta I'(a) + (1 - \beta)y < x = I'(a)$, we have that $y < x$. Hence, by the monotonicity of c, $d = \beta a + (1 - \beta)y^* \notin c(\{\beta a + (1 - \beta)y^*, \beta a + (1 - \beta)x^*\})$. Since $\beta a + (1 - \beta)x^* \in [a, x^*]$, it follows by Chernoff that $d \notin c(A \cup [a, x^*])$, as required.

It follows that $c(A \cup [a, x^*]) \subseteq [a, x^*]$. By Chernoff and Aizerman,[28] it follows that $c(A \cup [a, x^*]) = c([a, x^*]) = [a, x^*]$, where the second equality holds by property 2f of I', the constant linearity of c, and Chernoff.

□

Lemma IX.B.11. *I' represents c on O': for any $A \in \wp_{bdd}(O')$, $\arg\max_A I' \subseteq c(A)$.*

Proof. If $A \subseteq O$, the result follows from that fact that I' extends I and the fact that $(O, I) \in \mathcal{Z}$. So suppose that this is not the case, and consider $a \in \arg\max_A I'$. By Lemma IX.B.10, for each $P \in \mathcal{Q}$ with $A \cap P \neq \varnothing$, there exists $a'_P \in P$ with $I'(a'_P) = I'(a)$ and $c((A \cap P) \cup [a'_P, x^*_a]) = [a'_P, x^*_a]$. For such each P, let $A^a_P = (A \cap P) \cup [a'_P, x^*_a]$. By Chernoff and Aizerman,[29] $c(\bigcup_{P \text{ s.t } A \cap P \neq \varnothing} A^a_P) = c(\bigcup_{P \text{ s.t } A \cap P \neq \varnothing} [a'_P, x^*_a])$. More-

28. Note that, in the presence of Chernoff, Aizerman is equivalent to: $c(B) \subseteq A \subseteq B \Rightarrow c(B) = c(A)$. (See, for example, Moulin (1985).)

29. Chernoff and Aizerman imply that $c(\bigcup_{i \in I} A_i) = c(\bigcup_{i \in I} c(A_i))$. See for example Moulin (1985, Lemma 6), whose proof for the case of finite unions is straightforwardly extended to infinite unions.

over, $c(\bigcup_{P\ s.t\ A\cap P\neq\varnothing}[a'_P, x^*_a]) = \bigcup_{P\ s.t\ A\cap P\neq\varnothing}[a'_P, x^*_a]$, by property 2f of I', the constant linearity of c, and Chernoff. Hence, in particular, $a \in c(\bigcup_{P\ s.t\ A\cap P\neq\varnothing} A^a_P)$; since $A \subseteq \bigcup_{P\ s.t\ A\cap P\neq\varnothing} A^a_P$, it follows by Chernoff that $a \in c(A)$, as required.

□

So $(O', I') \in \mathcal{Z}$ such that $b \in O'$ and $(O, I) \leqslant (O', I')$. This concludes the proof of Lemma IX.B.4. □

Lemma IX.B.12. *$\mathcal{Z}$ is non-empty.*

Proof of Lemma IX.B.12. Define $C_{\hat{a}}$ as follows:

(IX.11)
$$C_{\hat{a}} = \left\{ y \in \Re \;\middle|\; \begin{array}{l} [\hat{a}, y^*] \subseteq c\left(X^{\hat{A}}_y\right) \,\&\, \\ \forall a' \in [\hat{a}, y^*],\ \forall z > 0, \forall \alpha \in (0,1), \\ \alpha\hat{a} + (1-\alpha)y^* \notin c\left(\{a' + z^*\} \cup X^{\hat{A}}_y\right) \end{array} \right\}$$

We first establish the non-emptyness of $C_{\hat{a}}$, by an argument related to, but not identical to, that used in the proof of Lemma IX.B.5. Consider $\left\{\sigma : [0,1) \to \Re \,\middle|\, \hat{a} \in c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in [0,1)\}\right)\right\}$. By the monotonicity of c and the fact that $\hat{a} \in c(\hat{A})$, this set is non-empty. Moreover, by the monotonicity and continuity of c, it has a maximal element; let σ be such an element. By the maximality of σ, $\hat{a} \in c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in [0,1]\}\right)$. For $\alpha \in [0,1)$ and $\epsilon > 0$, define $\sigma^{+\epsilon}_\alpha$ by: $\sigma^{+\epsilon}_\alpha(\alpha) = \sigma(\alpha) + \epsilon$ and $\sigma^{+\epsilon}_\alpha(\beta) = \sigma(\beta)$ for $\beta \neq \alpha$. By the maximality of σ, for any $\alpha \in [0,1)$ and $\epsilon > 0$, $\hat{a} \notin c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma^{+\epsilon}_\alpha(\alpha)^* \mid \alpha \in [0,1)\}\right)$. It follows by Aizerman that $\alpha\hat{a} + (1-\alpha)(\sigma(\alpha)+\epsilon)^* \in c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma^{+\epsilon}_\alpha(\alpha)^* \mid \alpha \in [0,1)\}\right)$. Since this holds for all $\epsilon > 0$ and $\alpha \in [0,1)$, it follows from continuity that $\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in [0,1)\} \subseteq c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in [0,1]\}\right)$. Let $y = \sigma(0)$; in particular, we have that $y^* \in c\left(\hat{A} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in [0,1)\}\right)$. It follows, by repeated applications of constant independence and Chernoff, that $\{\hat{\alpha}, y^*\} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \subseteq c\left(\{y^*\} \cup \bigcup_{k=0}^{n}(\delta^k\hat{A} + (1-\delta^k)y^*) \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\}\right)$. Taking the limit as $\delta \to 1$ and $n \to \infty$, it follows from continuity that $\{\hat{a}, y^*\} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \subseteq c\left(X^{\hat{A}}_y \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\}\right)$.

We now show that $\sigma(\alpha) = y$ for all $\alpha \in (0,1)$. Take any $\beta \in (0,1)$. Since $\beta X^{\hat{A}}_y + (1-\beta)y^* \subseteq X^{\hat{A}}_y$, $X^{\hat{A}}_y \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta(X^{\hat{A}}_y \cup \{\alpha\hat{a} + (1-$

$\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\}) + (1-\beta)y^*) = X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*)$. It follows from the constant independence of c that $\{\hat{a}, y^*\} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \subseteq c(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*))$. However, for any $z > 0$, by monotonicity, $\hat{a} \notin c\left(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup \{\hat{a} + z^*\}\right)$, so by constant linearity and Chernoff, $\beta\hat{a} + (1-\beta)y^* \notin c(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*) \cup \{\beta(\hat{a} + z^*) + (1-\beta)y^*\})$. It follows from Aizerman that $\beta(\hat{a} + z^*) + (1-\beta)y^* \in c(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*) \cup \{\beta(\hat{a} + z^*) + (1-\beta)y^*\})$; since this holds for every $z > 0$, we have, by continuity that $\beta\hat{a} + (1-\beta)y^* \in c(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*))$. However, as noted above, $\beta\hat{a} + (1-\beta)\sigma(\beta)^* \in c(X_y^{\hat{A}} \cup \{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} \cup (\beta\{\alpha\hat{a} + (1-\alpha)\sigma(\alpha)^* \mid \alpha \in (0,1)\} + (1-\beta)y^*))$; it follows from the monotonicity of c that $\sigma(b) = y$. Applying this to all $\beta \in (0,1)$, we have that the required conclusion. It follows in particular that $[\hat{a}, y^*] \subseteq c(X_y^{\hat{A}})$.

It remains to show that $\alpha\hat{a} + (1-\alpha)y^* \notin c\left(\{a' + z^*\} \cup X_y^{\hat{A}}\right)$ for all $\alpha \in (0,1)$, $a' \in [\hat{a}, y^*]$ and $z > 0$. We proceed by reductio: let $\alpha \in (0,1)$, $a' \in [\hat{a}, y^*]$ and $z > 0$ be such that $\alpha\hat{a} + (1-\alpha)y^* \in c\left(\{a' + z^*\} \cup X_y^{\hat{A}}\right)$ for all $\alpha \in (0,1)$. By definition $a' = \beta\hat{a} + (1-\beta)y^*$; we distinguish cases according to whether $\beta \leqslant \alpha$ or not. First suppose that $\beta \leqslant \alpha$. Chernoff implies that $\alpha\hat{a} + (1-\alpha)y^* \in c\left(\{\beta\hat{a} + (1-\beta)y^* + z^*\} \cup \{\alpha d + (1-\alpha)y^* \mid d \in \hat{A}\}\right)$. But since $\{\beta\hat{a} + (1-\beta)y^* + z^*\} \cup \{\alpha d + (1-\alpha)y^* \mid d \in \hat{A}\} = \alpha(\{(\frac{\beta}{\alpha}\hat{a} + (1-\frac{\beta}{\alpha})y^*) + \frac{1}{\alpha}z^*)\} \cup X_y^{\hat{A}}) + (1-\alpha)y^*$, it follows from constant linearity of c that $\hat{a} \in c(\{(\frac{\beta}{\alpha}\hat{a} + (1-\frac{\beta}{\alpha})y^*) + \frac{1}{\alpha}z^*)\} \cup X_y^{\hat{A}})$, contradicting the maximality of $\sigma(\frac{\beta}{\alpha}) = y$. Now suppose that $\beta > \alpha$. Take any $\gamma < \frac{\alpha}{\beta}$; by constant independence, $\alpha\hat{a} + (1-\alpha)y^* \in c\left(\{a' + z^*, \gamma(a' + z^*) + (1-\gamma)y^*\} \cup X_y^{\hat{A}}\right)$, and so, by Chernoff, $\alpha\hat{a} + (1-\alpha)y^* \in c\left(\{\gamma(a' + z^*) + (1-\gamma)y^*\} \cup X_y^{\hat{A}}\right)$. Since $\gamma(a' + z^*) + (1-\gamma)y^* = \gamma\beta\hat{a} + (1-\gamma\beta)y^* + \gamma z^*$, the conditions of the previous case are satisfied; the previous argument can thus be employed, yielding a contradiction. So $\alpha\hat{a} + (1-\alpha)y^* \notin c\left(\{a' + z^*\} \cup X_y^{\hat{A}}\right)$ for all $\alpha \in (0,1)$, $a' \in [\hat{a}, y^*]$ and $z > 0$, as required.

The construction of a functional I on $\overline{\hat{a}}$ proceeds in an analogous way to the proof of Lemma IX.B.4; the proofs that it satisfies the appropriate conditions are either trivial or follow the same reasoning as used in the proof of Lemma IX.B.4.

□

Bibliography

Al Najjar, N., Weinstein, J., and Al-Najjar, N. I. (2009). The Ambiguity Aversion Literature: A critical assessment. *Economics and Philosophy*, 25(Special Issue 03):249–284.

Aliprantis, C. D. and Border, K. C. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition.

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Billingsley, P. (2009). *Convergence of Probability Measures*. John Wiley & Sons.

Blackwell, D. (1953). Equivalent Comparisons of Experiments. *The Annals of Mathematical Statistics*, 24(2):265–272.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Rustichini, A. (2013). Niveloids and Their Extensions: Risk Measures in Small Domains.

Crès, H., Gilboa, I., and Vieille, N. (2011). Aggregation of multiple prior opinions. *Journal of Economic Theory*, 146:2563–2582.

Dekel, E., Lipman, B. L., and Rustichini, A. (2001). Representing Preferences with a Unique Subjective State Space. *Econometrica*, 69(4):891–934.

Dunford, N. and Schwartz, J. T. (1958). *Linear Operators. Part I: General Theory*. Interscience Publishers, New York.

Eichberger, J., Grant, S., and Kelsey, D. (2007). Updating Choquet beliefs. *Journal of Mathematical Economics*, 43(7–8):888–899.

Eichberger, J. and Kelsey, D. (1996). Uncertainty Aversion and Dynamic Consistency. *International Economic Review*, 37(3):625–640.

Epstein, L. G. and Le Breton, M. (1993). Dynamically Consistent Beliefs Must Be Bayesian. *Journal of Economic Theory*, 61(1):1–22.

Epstein, L. G., Marinacci, M., and Seo, K. (2007). Coarse contingencies and ambiguity. *Theoretical Economics*, 2:355–394.

Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Ghirardato, P. (2001). Coping with ignorance: unforeseen contingencies and non-additive uncertainty. *Economic Theory*, 17(2):247–276.

Ghirardato, P. (2002). Revisiting Savage in a conditional world. *Economic Theory*, 20(1):83–92.

Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *J. Econ. Theory*, 118(2):133–173.

Gilboa, I., Postlewaite, A., and Schmeidler, D. (2009). Is it always rational to satisfy {S}avage's axioms? *Economics and Philosophy*, pages 285–296.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Gollier, C. (2004). *The Economics Of Risk And Time*. MIT Press.

Hammond, P. J. (1988). Consequentialist foundations for expected utility theory. *Theory and Decision*, 25:25–78.

Hanany, E. and Klibanoff, P. (2007). Updating preferences with multiple priors. *Theoretical Economics*, 2:261–298.

Hill, B. (2012). Unanimity and the aggregation of multiple prior opinions.

Hilton, R. W. (1981). The Determinants of Information Value: Synthesizing Some General Results. *Management Science*, 27(1):57–64.

Karni, E. and Safra, Z. (1989). Dynamic Consistency, Revelations in Auctions and the Structure of Preferences. *The Review of Economic Studies*, 56(3):421–433.

Karni, E. and Safra, Z. (1990). Behaviorally consistent optimal stopping rules. *Journal of Economic Theory*, 51(2):391–402.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976.

Kreps, D. M. (1979). A Representation Theorem for ‘Preference for Flexibility’. *Econometrica*, 47:565–576.

Kreps, D. M. (1992). Static choice in the presence of unforeseen contingencies. In *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn*, pages 258–281.

LaValle, I. H. (1968). On Cash Equivalents and Information Evaluation in Decisions Under Uncertainty: Part I: Basic Theory. *Journal of the American Statistical Association*, 63(321):252–276.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Dynamic variational preferences. *Journal of Economic Theory*, 128(1):4–44.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668.

Marschak, J. and Miyasawa, K. (1968). Economic Comparability of Information Systems. *International Economic Review*, 9(2):137–174.

Marschak, J. and Radner, R. (1972). *Economic Theory of Teams*. Yale University Press, New Haven.

McClennen, E. F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press.

Moulin, H. (1985). Choice Functions Over a Finite Set: A Summary. *Social Choice Welfare*, 2:147–160.

Mukerji, S. (1997). Understanding the nonadditive probability decision model. *Economic Theory*, 9(1):23–46.

Sarin, R. and Wakker, P. P. (1998). Dynamic choice and nonexpected utility. *Journal of Risk and Uncertainty*, 17(2):87–120.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Sen, A. K. (1971). Choice Functions and Revealed Preference. *Rev. Econ. Stud.*, 38(3):307–317.

Siniscalchi, M. (2009). Two out of three ain't bad: a comment on "The Ambiguity Aversion Literature: A Critical Assessment”. *Economics and Philosophy*, 25(Special Issue 03):335–356.

Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics*, 6:379–421.

Strotz, R. H. (1955). Myopia and Inconsistency in Dynamic Utility Maximization. *Rev. Econ. Stud.*, 23(3):165–180.

Wakker, P. (1988). Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, I(July 1987):169–175.

X Uncertainty aversion, multi utility representations and state independence of utility

Abstract

This paper proposes and characterises a model of uncertainty averse preferences that can simultaneously accommodate three divergences from subjective expected utility: imprecision of beliefs (or ambiguity), imprecision of tastes (or multi utility), and state dependence of utility. Moreover, it characterises, in this context, a notion of state independence of utility borrowed from the literature on incomplete preferences. This notion is then shown to be basically inconsistent with the standard state-independence axiom, monotonicity, whenever tastes are imprecise. A new notion of state independence in the context of imprecise tastes, which is characterised by monotonicity, is proposed. [1]

Keywords: State independence of utility, imprecise tastes, uncertainty aversion, multi utility, multiple priors, state-dependent utility.

JEL classification: D81

1. This is a forthcoming HEC Paris Research Paper. It was writing during, and partially inspired by, the preparation of this thesis.

X.1 Introduction

Consider a decision maker who has been diagnosed with a serious yet rare degenerative disease, and is faced with the choice between not doing anything or paying to receive a new, expensive but not very well understood treatment. If he does nothing, there is a 40% chance of an uncomfortable but largely normal life (and lifespan) and a 60% chance of a significantly impeded lifestyle for the rest of his days. The chances of success of the treatment are between 25% and 75%; in case of success, it will be life as usual (with no discomfort); in case of failure, he will have a seriously degraded quality of life. This decision has a marked resemblance with Ellsberg's (1961) examples, insofar as the probabilities of various outcomes are known for one of the options, and they are not known precisely for the other. Borrowing a term from statistics, we will say that there may be *imprecision in beliefs*. [2] It is well-known that in such cases, the decision maker may exhibit non-neutral attitudes to uncertainty, such as uncertainty aversion. Moreover, given that few people know what it's like to live in a severely handicapped situation, this decision may also involve what could be called *imprecision in tastes*: the decision maker may not be able to compare the impairment resulting from failed treatment with the impeded quality of life resulting from inaction. Finally, as is well-known (Drèze, 1987; Arrow, 1974; Cook and Graham, 1977; Karni, 1983), such health-related decisions often involve state-dependent utility: one's utility for the money saved by not undergoing the treatment may depend on one's state of health.

Examples such as this, which we take as representative of an important class of real-life cases that are relevant in health economics, the theory of insurance and beyond, involve several distinct violations of the standard axioms of subjective expected utility (Savage, 1954; Anscombe and Aumann, 1963). [3] These violations have generally been studied in isolation: the literature on ambiguity, for example, almost universally assumes precise tastes and state independence of utility. The first aim of this paper is to propose and characterise a theory of decision that displays uncertainty aversion and is sufficiently rich to capture all of these

2. The term 'imprecise probabilities' is widely used in statistics and philosophy; see for example Walley (1991); Bradley (2014) and the references therein.

3. Explicitly, in the Anscombe-Aumann framework and in the context of complete preferences: violations of the independence axiom for imprecision of beliefs, of the restriction of independence to risky prospects for imprecision of tastes, and of the monotonicity axiom for state independence of utilities.

violations simultaneously. Such a theory will be applicable in the sorts of situations just described.

Our basic result will involve the representation of preferences by the following functional form, for acts (that is, state-contingent consequences) f:

$$\min_{U \in \mathcal{U}} \sum_{s \in S} U(s, f(s)) \tag{X.1}$$

where $\mathcal{U}$ is a (closed convex) set of real-valued functions on states and consequences, which we call *evaluations*.[4] Evaluations occur in the literature on state-dependent utility, where they have been recognised to provide a natural representation of preferences in cases where state-independence axioms are violated, which is sufficient for many economic applications. For example, they can be used to determine the choice made in the decision discussed above, as well as to perform comparative statics on, for example, the cost of the treatment (Hill, 2010). We provide a representation theorem showing that, in the presence of some mild technical conditions, the standard order, continuity and uncertainty aversion axioms – without monotonicity or any form of independence – are necessary and sufficient for preferences to be representable according to (X.1). It follows that (X.1) also generalises the vast majority of existing models of uncertainty averse preferences.

Despite its uses, representation (X.1) does not allow discussion of beliefs and tastes. For this, a separation is required. We shall focus on state-independent separations, of which the most notable example comes from standard subjective expected utility: there is a separation of the single evaluation in the singleton version of (X.1) into a (suitably unique) state-independent utility function and a probability measure, which is ensured by the state-independence axioms (monotonicity, in the Anscombe-Aumann framework). The main finding of this paper is that the issue of separation, and with it the meaning of state independence of utility, is considerably more complex when there is both uncertainty aversion and imprecision of tastes.

4. Some authors use the term 'state-dependent utilities' to refer to such functions. By contrast, we follow the literature on state-dependent utility (for example Karni et al. (1983); Karni (1993b,a); Karni and Schmeidler (1993); Karni and Mongin (2000); Karni (2011); Hill (2010)) in reserving the term 'state-dependent utility' for a representation involving a principled separation of beliefs and tastes: for example, a suitably unique probability and (state-dependent) utility.

Our first result is a characterisation of the following analogue, for uncertainty averse preferences, of a representation recently proposed by Galaabaatar and Karni (2013) in the context of incomplete preferences:

$$\min_{(u,p)\in\mathcal{U}^{si}} \sum_{s\in S} p(s)u(f(s)) \tag{X.2}$$

where $\mathcal{U}^{si}$ is a (closed convex) set of pairs of probabilities and state-independent utilities. Here there is separation evaluation-wise into a taste factor – the state-independent utility – and a belief factor – the probability. Since the state independence at issue holds at the level of each utility-probability pair, we shall refer to such representations as involving *multi state-independent utility*.

Surprisingly, this representation turns out to be in a sense inconsistent with the standard axiom for state independence, monotonicity. More precisely, multi state-independent utility and monotonicity taken together imply that, unless beliefs satisfy a particularly restrictive condition, preferences are represented by a single utility function. In other words, once there is imprecision in tastes, monotonicity – the standard axiom for state independence of utility – is incompatible with the notion of state independence embodied by (X.2) in all but a few special cases.

Investigating further, we show that, in the context of uncertainty averse preferences with imprecise tastes, monotonicity implies a particular representation of preferences, of which the following is a suggestive yet paradigmatic special case:

$$\min_{p\in\mathcal{C}} \sum_{s\in S} p(s) \min_{u\in\mathcal{U}^{ut}} u(f(s)) \tag{X.3}$$

where $\mathcal{C}$ is a (closed convex) set of probability measures on the state space and $\mathcal{U}^{ut}$ is a (closed convex) set of utility functions on the set of consequences. So monotonicity implies that the *set of utilities* involved is state independent, insofar as the same set is used in every state (in the second minimisation). We shall speak of *state-independent multi utility*, to emphasise that the state-independent separation applies at the level of sets, rather than individual utility-probability pairs.

These findings suggest that, once one leaves the realm of precise tastes, the single notion of state independence is replaced by two potentially relevant and basically incompatible notions. One, which is fast becoming the standard in studies of incomplete preferences, involves state independence at the level of each probability-utility pair involved in the representation. The other, which is implied by the most popular axiom for state independence, involves state independence at the level of the set of utilities. Of course, the difference between them may have consequences for behaviour, and in applications.

The paper is organised as follows. The framework is set out in Section X.2. The basic result is stated and discussed in Section X.3, which also contains a corollary for preferences over risky prospects. Section X.4 contains the analysis of the issue of state independence: Section X.4.1 focusses on multi state-independent utility and its incompatibility with monotonicity; Section X.4.2 considers state-independent multi utility. Related literature is discussed in Section X.5. Proofs are contained in an Appendix.

X.2 Framework

We use the standard Anscombe-Aumann (1963) framework, as adapted by Fishburn (1970). Let S be a finite set of states; Δ is the set of probability measures over S. Let X be a finite set of outcomes. *Consequences* are *lotteries* over X: that is, probability measures over X. $\Delta(X)$ is the set of consequences; of course $\Delta(X) \subset \Re^X$. An *act* is a function from states to consequences; let $\mathcal{A} = \Delta(X)^S$ be the set of acts. $\mathcal{A}$ is a mixture set with the mixture relation defined pointwise: for f, h in $\mathcal{A}$ and $\alpha \in [0, 1]$, the mixture $\alpha f + (1-\alpha)h$ is defined by $(\alpha f + (1-\alpha)h)(s, x) = \alpha f(s, x) + (1-\alpha)h(s, x)$. We write $f_\alpha h$ as short for $\alpha f + (1-\alpha)h$. For state s and acts f and g, the act $f_s g_{s^c}$ is defined as follows: $f_s g_{s^c}(s) = f(s)$ and $f_s g_{s^c}(t) = g(t)$ for all $t \neq s$. With slight abuse of notation, a constant act taking consequence c for every state will be denoted c and the set of constant acts will be denoted $\Delta(X)$.

An *evaluation* is a function $U : S \times X \to \Re$. We assume $\Re^{S\times X}$ as well as $\mathcal{A}$ to be endowed with the Euclidean norm $|\,|$. For evaluation $U \in \Re^{S\times X}$, consequence $c \in \Delta(X)$ and state $s \in S$, we let $U(s, c) = \sum_{x\in X} U(s, x)c(x)$. Evaluations $U, U' \in \Re^{S\times X}$ are *cardinally equivalent* if there exists $b \in \Re^S$ with $\sum_{s\in S} b(s) = 0$ such that $U' = U + b$.[5] $U' \in \Re^{S\times X}$ is a *positive affine transformation* of U if there exists $a \in \Re_{>0}$ and $b \in \Re$ such that $U' = aU + b$. An evaluation representing preferences according to (X.1) with a singleton $\mathcal{U}$ is unique up to cardinal equivalence and positive affine transformation (see

5. Cardinally equivalent evaluations are such that $\sum_{s\in S} U(s, f(s)) = \sum_{s\in S} U'(s, f(s))$ for all $f \in \mathcal{A}$.

for example Karni et al. (1983)). We extend these notions to sets of evaluations as follows. For $\mathcal{U}, \mathcal{U}' \subseteq \Re^{S\times X}$, $\mathcal{U}'$ is cardinally equivalent to $\mathcal{U}$ if, for each $U' \in \mathcal{U}'$, there exists $U \in \mathcal{U}$ such that U and U' are cardinally equivalent, and similarly for each $U \in \mathcal{U}$. $\mathcal{U}'$ is a positive affine transformation of $\mathcal{U}$ if there exist $a \in \Re_{>0}$ and $b \in \Re$, such that $\mathcal{U}' = a\mathcal{U} + b = \{aU + b \mid U \in \mathcal{U}\}$. We adopt analogous notation and terminology for state-independent utilities (ie. functions in $\Re^X$). Fix an arbitrary $\hat{x} \in X$, and let $\mathcal{N}^1 = \{U \in \Re^{S\times X} \mid \forall s \in S,\ U(s,\hat{x}) = 0,\ \ |U| = 1\}$. Up to cardinal equivalence and positive affine transformation, $\mathcal{N}^1$ contains a representative of each evaluation: it can thus be thought of as a set of appropriately normalised evaluations.

The binary relation $\preceq$ on $\mathcal{A}$ represents the decision maker's preferences over acts. The symmetric and asymmetric parts of $\preceq$, $\sim$ and $\prec$, are defined in the standard way. A functional $V : \mathcal{A} \to \Re$ *represents* $\preceq$ if and only if, for all $f, g \in \mathcal{A}$, $f \preceq g$ iff $V(f) \leqslant V(g)$.

An evaluation U is *constant* if $\sum_{s\in S} U(s, f(s)) = \sum_{s\in S} U(s, g(s))$ for all $f, g \in \mathcal{A}$. A set $\mathcal{U} \subseteq \Re^{S\times X}$ is *non-trivial* if it contains at least one non-constant evaluation. Let $\mathcal{F} \subseteq 2^{\Re^{S\times X}}$ be a family of closed convex sets of evaluations (for example, $\mathcal{F}$ could contain all closed convex sets). For any $\mathcal{U} \in \mathcal{F}$ representing $\preceq$ according to (X.1), let $\mathcal{U}(\mathcal{A}) = \{\min_{U\in\mathcal{U}} \sum_{s\in S} U(s, f(s)) \mid f \in \mathcal{A}\}$. Since $\mathcal{A}$ is closed, and the minimal functional and the U are continuous, this is a closed set. For any $\mathcal{U} \in \mathcal{F}$ representing $\preceq$ according to (X.1), we shall say that $\mathcal{U}$ is a *tight* closed convex representation of $\preceq$ from the family $\mathcal{F}$ if there exists no proper subset in the family, $\mathcal{U}' \subsetneq \mathcal{U}$ with $\mathcal{U}' \in \mathcal{F}$, representing $\preceq$ according to (X.1). A tight set is as small as a representation can be, in the sense that there are no extraneous evaluations. $\mathcal{U}$ is a *minimal* tight closed convex set from $\mathcal{F}$ representing $\preceq$ if for every tight closed convex set $\mathcal{U}' \in \mathcal{F}$ representing $\preceq$ with $\mathcal{U}'(\mathcal{A}) = \mathcal{U}(\mathcal{A})$, and for each $f \in \mathcal{A}$ and $U' \in \mathcal{U}'$ there exists $U \in \mathcal{U}$ such that $\sum_{s\in S} U'(s, f(s)) \geqslant \sum_{s\in S} U(s, f(s))$. A minimal representing set contains evaluations taking the minimal values possible, compared to other representing sets that are suitably normalised (so as to take the same range of values). Explicit mention of the family $\mathcal{F}$ will be omitted where it is evident from the context. Again, analogous notions are employed for sets of state-independent utilities (ie. functions in $\Re^X$).

X.3 The base result

X.3.1 Axioms

Consider the following axioms.

Axiom X-A1 (Weak order)**.** For all $f, g, h \in \mathcal{A}$: if $f \preceq g$ and $g \preceq h$, then $f \preceq h$; and $f \preceq g$ or $g \preceq f$.

Axiom X-A2 (Non-degeneracy)**.** There exists $f, g \in \mathcal{A}$ such that $f \succ g$.

Axiom X-A3 (Continuity)**.** For all sequences $f_n, g_n \in \mathcal{A}$, $f_n \to f$ and $g_n \to g$, if $f_n \succeq g_n$ for all n, then $f \succeq g$.

Axiom X-A4 (Uncertainty aversion)**.** For all $f, g \in \mathcal{A}$ and $\alpha \in (0, 1)$, if $f \sim g$ then $f_\alpha g \succeq g$.

Axiom X-A5 (Local non-satiation)**.** For every $f \in \mathcal{A}$, either $f \succeq g$ for all $g \in \mathcal{A}$ or there exists $g \in \mathcal{A}$ such that $g_\alpha f \succ f$ for all $\alpha \in (0, 1]$.

X-A1–X-A4 require little comment: X-A1 and X-A2 are totally standard, X-A3 is a standard topological continuity condition (equivalence with the more common mixture continuity axiom does not evidently hold here, given the absence of any independence axiom),[6] and X-A4 is the standard uncertainty aversion axiom from Schmeidler (1989); Gilboa and Schmeidler (1989). Local non-satiation, X-A5, is a version of a standard condition in consumer theory (for example, (Mas-Colell et al., 1995, Ch 3)), which ensures that all indifference curves are 'thin'. It is a consequence of monotonicity and independence, which are not generally assumed here.

One further axiom shall be of interest, for which the following notion will prove useful. For acts $f, g, h \in \mathcal{A}$ with $g \prec f \prec h$, f is a *potential fifty-fifty mixture* of g and h if there exist $g', h' \in \mathcal{A}$ with $g' \sim g$, $h' \sim h$ and $h'_{\frac{1}{2}} g' \sim f$. The potential fifty-fifty mixtures of g and h are those acts that are indifferent to a fifty-fifty mixture of – or if you will, to a fair coin toss yielding – two acts which are indifferent to g and h respectively. In the presence of independence, the set of potential fifty-fifty mixtures is the indifference class of the

6. Cerreia-Vioglio (2009) shows that mixture continuity implies upper semi-continuity in our framework; however, the result used in the proof of Theorem X.1 (by Kannai (1977)) requires full continuity.

fifty-fifty mixture of g and h; to this extent, the notion of potential fifty-fifty mixture can be thought of as an analogue of the standard notion of fifty-fifty mixture for cases where the axiom is not satisfied. Moreover, f is a *conservative potential fifty-fifty mixture* of g and h if it is a potential fifty-fifty mixture and for any other potential fifty-fifty mixture, f', of these acts $f' \geq f$. Conservative potential fifty-fifty mixtures are the lowest acts, in the $\preceq$ order, that are potential fifty-fifty mixtures.

Now consider the following axiom.

Axiom X-A6 (Coherent calibration)**.** For any $f, g, h \in \mathcal{A}$ with $g \prec f \prec h$, there exists $\epsilon > 0$ such that, for every sequence $(f_i)_{0 \leqslant i \leqslant n}$ with $f_0 = g$, $f_n = h$, $f_i \succ f_{i-1}$ for all $1 \leqslant i \leqslant n$ and such that f_i is a conservative potential fifty-fifty mixture of f_{i-1} and f_{i+1} for all $1 \leqslant i \leqslant n-1$, $\frac{\max_{f_i \preceq f} i}{n} \leqslant 1 - \epsilon$.

A sequence of acts, each consecutive triple of which involves a conservative potential fifty-fifty mixture, can be thought of as forming a 'ruler' with equally-spaced gaps (in terms of preferences) between the elements of the sequence. (As such, there is conceptual analogy with the notion of standard sequence common in measurement theory (Krantz et al., 1971).) Coherent calibration, X-A6, demands that for any act strictly between two acts g and h, every such ruler gives a measurement of the position of f that is strictly bounded away from the distance between g and h. The intuition is obvious: if $f \prec h$, then f must be strictly 'closer' to g, according to any such ruler, than h. One could consider this axiom as guaranteeing that sequences of conservative potential fifty-fifty mixtures constitute reasonable 'rulers' to calibrate preferences. (Note that this is automatically the case in the presence of independence.) As will be noted below, the vast majority of uncertainty averse decision models proposed in the literature satisfy X-A6.

X.3.2 Representation Theorem

The following is the first basic result of the paper.

Theorem X.1. *Let $\preceq$ be a preference relation on $\mathcal{A}$. The following are equivalent:*

(i) *$\preceq$ satisfies X-A1–X-A6,*

(ii) *there exists a non-trivial minimal tight closed convex set of evaluations $\mathcal{U} \subseteq \Re^{S \times X}$ such that $\preceq$ is represented by:*

$$V(f) = \min_{U \in \mathcal{U}} \sum_{s \in S} U(s, f(s)) \tag{X.1}$$

Moreover, for any minimal tight closed convex set $\mathcal{U}' \subseteq \Re^{S \times X}$ representing $\preceq$ according to (X.1) there exist $\mathcal{U}^\star \subseteq \Re^{S \times X}$, $a \in \Re_{>0}$ and $b \in \Re$ such that $\mathcal{U}$ is cardinally equivalent to $\mathcal{U}^\star$ and $\mathcal{U}' = a\mathcal{U}^\star + b$.

Hence the axioms given above – which, bar two largely technical assumptions, boil down to standard ordering, continuity and uncertainty aversion – characterise the multi evaluation representation of preferences mentioned in the Introduction. Since no monotonicity or independence assumptions are involved (not even independence on risky prospects), this representation is sufficiently rich to capture simultaneous violations of these axioms. In particular, it can accommodate state dependence of utility, imprecision of tastes and imprecision of beliefs, when the preferences exhibit aversion to uncertainty.

Indeed, representation (X.1) can be seen as a natural development of existing research focussing on some of these violations in isolation from the others. It is a natural extension of the single evaluation representation (where $\mathcal{U}$ is a singleton), which is obtained when one removes the state-independence axioms from the standard axiomatisation of subjective expected utility.[7] Moreover, it is a natural analogue, for uncertainty averse preferences, of recent representations of incomplete preferences obtained by Seidenfeld et al. (1995); Nau (2006); Galaabaatar and Karni (2013), which demand preference only when the expression in (X.1) is higher for all evaluations in $\mathcal{U}$. Such representations can accommodate imprecision in tastes and state dependence of utility, though not uncertainty aversion of preferences. Finally, representation (X.1) can also be thought of as a generalisation of the popular Maxmin Expected Utility model (Gilboa and Schmeidler, 1989) in the ambiguity literature, insofar as it involves minimisation over sets of evaluations as opposed to minimisation over sets of probability measures (with a single state-independent utility function). In fact, despite its prima facie resemblance to this model, the use of the minimum over a set of evaluations rather than something more refined is not restrictive: representation (X.1) generalises a large class of ambiguity models.

7. This is a straightforward consequence of the von Neumann-Morgenstern theorem in the Anscombe-Aumann framework (Karni et al., 1983). Wakker and Zank (1999); Hill (2010) propose analogous results in the Savage framework.

This is clear from the proof of the Theorem (see Appendix X.A.1), where it is shown that the axioms are equivalent to the representation of preferences by a continuous concave functional on $\mathcal{A}$. Since the vast majority of existing uncertainty averse decision theories involve continuous concave functionals – including the uncertainty averse Choquet preferences (Schmeidler, 1989), maxmin EU preferences (Gilboa and Schmeidler, 1989), variational preferences (Maccheroni et al., 2006), confidence preferences (Chateauneuf and Faro, 2009), and uncertainty averse smooth preferences (Klibanoff et al., 2005) – they are special cases of representation (X.1). This also indicates how mild the only new axiom, Coherent calibration (X-A6), is. To the knowledge of the author, the only uncertainty averse theory that does not fall under representation (X.1) is the general case of the uncertainty averse preference representation given in Cerreia-Vioglio et al. (2011b).

X.3.3 An Alternative Formulation

For future reference, it is useful to note that, by standard duality (Rockafellar, 1970), representation (X.1) is equivalent to the representation of $\preceq$ by the functional

$$V(f) = \min_{U\in\mathcal{N}^1,(a,b)\in\alpha(U)} \left(\sum_{s\in S} a.U(s,f(s)) + b(s)\right) \tag{X.4}$$

where $\alpha : \mathcal{N}^1 \to 2^{\Re_{>0}\times\Re}$ is a function with the following properties: [8]

- non-triviality: there exists non-constant $U \in \mathcal{N}^1$ such that $\alpha(U) \neq \varnothing$;
- convexity: if $(a,b) \in \alpha(U)$, $(a',b') \in \alpha(U')$ then for all $\lambda \in [0,1]$, $(\lambda a + (1-\lambda)a', \lambda b + (1-\lambda)b') \in \alpha(\lambda U + (1-\lambda)U')$;
- upper hemicontinuity: if $(a_n,b_n) \in \alpha(U_n)$ with $U_n \to U$ and $(a_n,b_n) \to (a,b)$, then $(a,b) \in \alpha(U)$;
- tightness: there exists no convex upper hemicontinuous α' representing $\preceq$ according to (X.4) with $\alpha'(U) \subseteq \alpha(U)$ for all $U \in \mathcal{N}^1$ where the inclusion is proper for at least one $U \in \mathcal{N}^1$;
- minimality: for any other tight convex upper hemicontinuous α' representing $\preceq$ according to (X.4) such that $\{\min_{U\in\mathcal{N}^1,(a,b)\in\alpha'(U)} \left(\sum_{s\in S} a.U(s,f(s)) + b(s)\right) | f \in$

8. Note that α is not a correspondence: $\alpha(U)$ may be empty for some $U \in \mathcal{N}^1$. We adopt the convention that when $\alpha(U)$ is empty, U is not involved in the minimisation.

$\mathcal{A}\} = \{\min_{U \in \mathcal{N}^1, (a,b) \in \alpha(U)} \left(\sum_{s \in S} a.U(s, f(s)) + b(s)\right) | f \in \mathcal{A}\}$, and for all $f \in \mathcal{A}$, $U' \in \mathcal{N}^1$ and $(a', b') \in \alpha'(U')$, there exists $U \in \mathcal{N}^1$ and $(a, b) \in \alpha(U)$ with $\sum_{s \in S} a.U(s, f(s)) + b(s) \leqslant \sum_{s \in S} a'.U'(s, f(s)) + b'(s)$.

In fact, we have the following direct translation between representations (X.1) and (X.4): [9]

$$\alpha(U) = \{(a, b) \mid a > 0, b \in \Re,\ aU + b \in \mathcal{U}^\star\} \tag{X.5}$$

$$\mathcal{U}^\star = \left\{aU + b \mid U \in \mathcal{N}^1,\ (a, b) \in \alpha(U)\right\} \tag{X.6}$$

where $\mathcal{U}^\star$ is the set of evaluations cardinally equivalent to $\mathcal{U}$ such that $U(s, \hat{x}) = U(t, \hat{x})$ for all $s, t \in S$ and $U \in \mathcal{U}^\star$. [10]

Instead of involving a set of evaluations, representation (X.4) works with suitably normalised evaluations – elements of $\mathcal{N}^1$. As noted in Section X.2, the elements of $\mathcal{N}^1$ suffice to represent preferences according to the version of (X.1) with a singleton set: any preference relation represented by an evaluation in this way can be represented by a member of $\mathcal{N}^1$. When several evaluations are involved in the representation, the function α determines how they are 'calibrated' for the decision maker (the a and b determine where the zero and unit are). Given the use of the maxmin rule, this in turn determines the 'weight' given to the different evaluations in the assessment of an act. [11]

X.3.4 Risk preferences

It is useful to remark that, applied in the special case of a state space containing a single state, Theorem X.1 yields the following axiomatisation of preferences under risk.

Corollary X.1. *Let $\preceq$ be a preference relation on $\Delta(X)$. $\preceq$ satisfies X-A1–X-A6 if and only if there exists a non-trivial minimal tight closed convex set of utility functions $\hat{\mathcal{U}} \subseteq \Re^X$ such that $\preceq$ is represented by:*

9. The equivalence of the representations and the properties of α follow directly from this translation.
10. Recall that $\hat{x}$ is an arbitrary outcome used in the normalisation defining $\mathcal{N}^1$; see Section X.2.
11. This can be formulated alternatively using a set of evaluations-as-a-group analogy: one could consider the elements of $\mathcal{N}^1$ as representing the different possible preferences of members of a group and the α as setting the interpersonal comparison of evaluations among the group members.

$$V(c) = \min_{u \in \hat{\mathcal{U}}} u(c) \tag{X.7}$$

Moreover, for any minimal tight closed convex $\hat{\mathcal{U}}' \subseteq \Re^X$ *representing* $\preceq$ *according to (X.7) there exists* $a \in \Re_{>0}$ *and* $b \in \Re$ *such that* $\hat{\mathcal{U}}' = a\hat{\mathcal{U}} + b$*.*

Moreover, it follows that, for each $\mathcal{U}$ representing a preference relation $\preceq$ over $\mathcal{A}$ (with a potentially non-singleton state space), one can define a unique 'restriction' of $\mathcal{U}$ which represents the restriction of $\preceq$ to constant acts $\Delta(X)$.

Proposition X.1. *Let* $\preceq$ *be a preference relation on* $\mathcal{A}$*. Suppose that* $\preceq$ *satisfies X-A1–X-A6, and is represented by a non-trivial minimal tight closed convex* $\mathcal{U} \subseteq \Re^{S\times X}$ *according to (X.1). Then there exists a unique non-trivial tight closed convex* $\mathcal{U}_{\Delta(X)} \subseteq \Re^X$ *representing the restriction of* $\preceq$ *to constant acts* $\Delta(X)$ *according to (X.7), such that for each* $u \in \mathcal{U}_{\Delta(X)}$*, there exists* $U \in \mathcal{U}$ *with* $u(x) = \sum_{s\in S} U(s,x)$ *for all* $x \in X$*.*

Henceforth, for $\mathcal{U}$ representing $\preceq$, we call the set $\mathcal{U}_{\Delta(X)}$ identified in this Proposition the *restriction of* $\mathcal{U}$ *to constant acts* (and continue to use the notation $\mathcal{U}_{\Delta(X)}$).

X.4 A Tale of Two State Independences

Representation (X.1) does not provide the basis for consideration and discussion of beliefs and tastes. For this, a representation involving utilities (the standard representations of tastes) and probabilities (which are often interpreted as beliefs) is required. Representation (X.1) does however provide a springboard for studying such representations: the task is to provide a principled separation of the set of evaluations $\mathcal{U}$ into utilities and probabilities. [12] We shall focus here on representations involving state independence of utility.

Galaabaatar and Karni (2013) provide inspiration. Working with incomplete preferences, they consider a 'dominance' condition that, under their other assumptions, is stronger than the monotonicity axiom standard in the Anscombe-Aumann framework. They

12. As has been noted in the literature on state-dependent utility, any evaluation can be separated into a (state-dependent) utility and probability in many ways. The challenge is to provide a suitably unique separation. Demanding state-independence of the utility is a possible way of doing so, though it is not the only one.

show that, in the context of a representation that may be thought of as an analogue of (X.1) for incomplete preferences, this condition is necessary and sufficient for preferences to be representable as follows: for $f, g \in \mathcal{A}$, $f \prec g$ if and only if

$$\sum_{s \in S} p(s)u(f(s)) < \sum_{s \in S} p(s)u(g(s)) \qquad \text{for all } (p,u) \in \mathcal{U}^{si} \tag{X.8}$$

where $\mathcal{U}^{si}$ is a set of pairs of probability measures over S and state-independent utilities over X. In other words, they study a notion of state independence that involves the separation of *each evaluation* into a state-independent utility and a probability, yielding a set of state-independent utility-probability pairs. This representation involves what we called multi state-independent utility.

In this section, we first characterise uncertainty averse preferences involving multi state-independent utility, and then go on to consider in depth the relationship with and consequences of monotonicity.

X.4.1 Multi state-independent utility

X.4.1.1 Basic state-independent representation

We begin by introducing a new axiom. To formulate it, we shall introduce the notion of δ-shift. Firstly, let $\mathcal{ZUR} = \{e \in \Re^X | \sum_{x \in X} e(x) = 0, \ |e| = 1\}$. For every $c \in \Delta(X)$, $\delta \in (0,1)$ and $e \in \mathcal{ZUR}$, the *δ-shift of c by e* is $c + \delta e$. If $c + \delta e \in \Delta(X)$ we shall say that that δ-shift exists. A δ-shift from c can be thought of as an act obtained by the addition of a small zero-sum risk, δe, to c. ($\mathcal{ZUR}$ is the set of zero-sum unit risks.) We denote the δ-shift of c by e, if it exists, by $c^{e,\delta}$.

A state $s \in S$ is *locally null* at $f \in \mathcal{A}$ if, for each $e \in \mathcal{ZUR}$, there exists $\epsilon > 0$ such that for all $\delta \in (0, \epsilon]$, $f(s)^{e,\delta}$ exists and $f(s)^{e,\delta}_s f_{s^c} \sim f$. For every $s \in S$, let $Null(s) = cl\left(\{f \in \mathcal{A} \mid s \text{ locally null at } f\}\right)$.[13] $s \in S$ is *locally nonnull* at $f \in \mathcal{A}$ if $f \notin Null(s)$.

Axiom X-A7 (Local state independence)**.** For all $f \in \mathcal{A}$, $s, t \in S$ locally nonnull at f, and all $e \in \mathcal{ZUR}$ such that $f(s)^{e,\delta}$ and $f(t)^{e,\delta}$ exist for some $\delta > 0$, there exists $\epsilon > 0$ such that $f(s)^{e,\delta}_s f_{s^c} \succeq f$ for all $\delta \leqslant \epsilon$ if and only if $f(t)^{e,\delta}_t f_{t^c} \succeq f$ for all $\delta \leqslant \epsilon$.

13. For a set X, $cl(X)$ is the closure of X.

This axiom has the spirit of a state-independence condition, to the extent that it demands that adding a small zero-sum risk at a nonnull state of an act will have the same effect on preferences, irrespective of the nonnull state where it is added. It is local, because this only holds for sufficiently small perturbations to the consequences obtained in the relevant states.

It turns out that this axiom is precisely what is required for a refinement of representation (X.1) where each evaluation is decomposed (uniquely) into a state-independent utility function and a probability measure.

Theorem X.2. *Let $\preceq$ be a preference relation on $\mathcal{A}$. The following are equivalent:*

- *(i)* *$\preceq$ satisfies X-A1–X-A6 and X-A7,*
- *(ii)* *there exists a non-trivial minimal tight closed convex set of utility-probability pairs $\mathcal{U}^{si} \subseteq \Re^X \times \Delta$ such that $\preceq$ is represented by:*

$$V(f) = \min_{(u,p)\in\mathcal{U}^{si}} \sum_{s\in S} p(s)u(f(s)) \tag{X.2}$$

Moreover, for any minimal tight closed convex $\mathcal{U}^{si\prime} \subseteq \Re^X \times \Delta$ representing $\preceq$ according to (X.2), there exist $a \in \Re_{>0}$ and $b \in \Re$ such that $\mathcal{U}^{si\prime} = a\mathcal{U}^{si} + b$. [14]

This representation, as well as the uniqueness obtained, can be thought of as the analogue of Galaabaatar and Karni's representation (X.8) for uncertainty averse preferences. Like theirs, it involves a state-independent separation of beliefs and tastes (probabilities and utilities) evaluation-by-evaluation. The notion of state independence at issue here applies at the level of each utility-probability pair; for this reason, we call (X.2) a *multi state-independent utility (multi prior) representation*.

X.4.1.2 Set factorisation

Representation (X.8) involves a set of state-independent utility-probability pairs, but it does not guarantee that this set can be factorised into a set of state-independent utilities and a set of probabilities. Galaabaatar and Karni (2013) go on to provide a characterisation of

14. It is understood in this result and those below that minimality and tightness are defined as stated in Section X.2, with the obvious family $\mathcal{F}$; in this case $\mathcal{F}$ is the set of all closed convex subsets of $\Re^X \times \Delta$.

what they dub a 'complete separation of beliefs and tastes'. It is possible to provide an analogous extension of Theorem X.2 in the case of uncertainty averse preferences.

To this end, we formulate a strong version of the Belief consistency axiom proposed by Galaabaatar and Karni (2013). Like them, we shall make use of the following notion, introduced by Ok et al. (2012). For an act $f \in \mathcal{A}$ and a probability measure $p \in \Delta$, let $f^p \in \Delta(X)$ be defined by: $f^p(x) = \sum_{s \in S} p(s) f(s,x)$.[15] Let $\mathcal{P} = \{p \in \Delta \mid \forall f \in \mathcal{A},\ f^p \succeq f\}$.

Axiom X-A8 (Strong belief consistency)**.** For all $f, g \in \mathcal{A}$, $f \succeq g$ if and only if for all $p \in \mathcal{P}$, there exists $q \in \mathcal{P}$ such that $f^p \succeq g^q$.

Strong belief consistency can be interpreted in a similar way to the Belief consistency axiom proposed by Galaabaatar and Karni (2013). They point out that an act in the Anscombe-Aumann framework can be thought of as a tacit compound lottery, with as probabilities in the first stage the subjective probabilities that implicitly govern behaviour. They go on to interpret f^p as the reduction of this compound lottery (when the subjective probability is p). Applying this to sets of probability measures, one comes naturally to the idea that preferences over acts should correspond to preferences over appropriate reductions, where the reductions are taken with respect to the 'probability distributions that are consistent with preferences', as they call them (in the current case, the set $\mathcal{P}$).

This axiom ensures the desired factorisation into a set of state-independent utility functions and a set of probability measures, as the following result shows.

Theorem X.3. *Let $\preceq$ be a preference relation on $\mathcal{A}$. The following are equivalent:*

(i) *$\preceq$ satisfies X-A1–X-A6, X-A7 and X-A8,*

(ii) *there exists a non-trivial minimal tight pair of closed convex sets of utility functions and probability measures, $\mathcal{U}^{ut} \subseteq \Re^X$ and $\mathcal{C} \subseteq \Delta$, such that $\preceq$ is represented by:*

$$V(f) = \min_{u \in \mathcal{U}^{ut},\ p \in \mathcal{C}} \sum_{s \in S} p(s) u(f(s)) \tag{X.9}$$

Moreover, for any minimal tight pair of closed convex sets $\mathcal{U}^{ut\prime} \subseteq \Re^X$ and $\mathcal{C}' \subseteq \Delta$ representing $\preceq$, $\mathcal{C}' = \mathcal{C}$ and there exist $a \in \Re_{>0}$ and $b \in \Re$ such that $\mathcal{U}^{ut\prime} = a\mathcal{U}^{ut} + b$.[16]

15. The reader is referred to the cited papers for extended discussion of the interpretation of this notion.

16. Minimality and tightness are defined as stated in Section X.2, with the obvious family $\mathcal{F}$; in this case $\mathcal{F}$ is the set of all closed convex subsets of $\Re^X \times \Delta$ that can be factorised into the product of a subset of $\Re^X$ and a subset of Δ.

X.4.1.3 Multi state-independent utility and Monotonicity

Despite its adequacy for characterising representations involving multi state-independent utility, local state independence (X-A7) is far from standard. The standard state-independence condition in the Anscombe-Aumann framework is the following monotonicity axiom. [17]

Axiom X-A9 (Monotonicity)**.** For all $f, g \in \mathcal{A}$, if $f(s) \succeq g(s)$ for all $s \in S$, then $f \succeq g$.

The vast majority, if not all, ambiguity theories incorporate this axiom. Moreover, as noted in the Introduction, the vast majority also assume precise tastes. We shall say that a preference relation $\preceq$ represented by a minimal tight closed convex $\mathcal{U}$ according to (X.1) involves:

— *precise tastes* if $\mathcal{U}_{\Delta(X)}$ is a singleton;

— *rudimentary beliefs* if there exists $E \subseteq S$ and concave real-valued functions $\phi_s : \mathcal{U}(\mathcal{A}) \to \Re$ for all $s \in E$ such that $\preceq$ is represented by:

$$V(f) = \min_{s \in E} \phi_s \left(\min_{u \in \mathcal{U}_{\Delta(X)}} u((f(s)) \right) \tag{X.10}$$

Preferences involve precise tastes if there is a single utility function (recall from Section X.3.4 that $\mathcal{U}_{\Delta(X)}$ is the restriction of $\mathcal{U}$ to constant acts). They involve rudimentary beliefs if there is no probabilistic dimension to beliefs: they are summed up by a set of states, over which the decision maker applies a maxmin rule.

For $c, c' \in \Delta(X)$, $c \prec c'$, let $\preceq_{[c,c']}$ be the restriction of $\preceq$ to $\{f \in \mathcal{A} \mid \forall s \in S,\ c' \preceq f(s) \preceq c\}$. We call $\preceq_{[c,c']}$ a *section* of the preferences. If $\preceq$ can represented according to (X.1), then any section can also be represented this way. A preference $\preceq$ representable according to (X.1) exhibits *section-wise precise tastes-rudimentary beliefs* if there exists a sequence $c_0, \ldots, c_n \in \Delta(X)$ with $c_0 \preceq f \preceq c_n$ for all $f \in \mathcal{A}$ and $c_i \prec c_{i+1}$ for each $i < n$, such that for each $i < n$, $\preceq_{[c_i, c_{i+1}]}$ involves precise tastes or rudimentary beliefs.

The following proposition sheds some light on the relationship between multi state-independent utility and monotonicity.

Proposition X.2. *Let $\preceq$ be a preference relation on $\mathcal{A}$ represented according to (X.2). If $\preceq$ satisfies X-A9, then it exhibits section-wise precise tastes-rudimentary beliefs.*

17. In the presence of complete transitive preferences, monotonicity is implied by the following alternative condition for state independence: for every $f, g \in \mathcal{A}$, $c, d \in \Delta(X)$ and $s, t \in S$, $c_s f_{s^c} \preceq d_s f_{s^c}$ if and only if $c_t g_{t^c} \preceq d_t g_{t^c}$. The main results stated below thus continue to hold for this state-independence condition.

This result says that whenever a multi state-independent utility representation satisfies monotonicity, it basically involves either a single utility function or non-probabilistic beliefs at every point in the preference order. It can be interpreted as an impossibility result: in the presence of all but the most simple beliefs, imprecision of tastes, multi state-independent utility and the monotonicity axiom are incompatible, when preferences are uncertainty averse. [18] If one takes the idea of imprecise tastes seriously, there is thus a tension between the prevalent notion of state independence for multi utility multi prior representations – according to which each utility in the representing set is state independent – and the standard axiom for state independence – namely monotonicity.

To gain some intuition for the result, consider the following simple example, to which we shall return below. Suppose that there are two states, s and t and two outcomes c and d. Let u_1, u_2 be state-independent utilities with $u_1(c) = 0$, $u_1(d) = 1$, $u_2(c) = 1$, $u_2(d) = 0$ and p be a probability with $p(s) = \frac{1}{4}$. Consider a multi state-independent utility representation (X.2) by the set $\mathcal{U}^{si}$, which is the convex closure of the following utility-probability pairs: (u_1, p), (u_2, p). It is evident that according to this representation $c \sim d$, but $c_s d_{s^c} \succ c$, contradicting monotonicity.

Proposition X.2 suggests that, when tastes are imprecise, one can retain either multi state-independent utility or monotonicity, but not both. We have already considered one of the options; it is time to investigate the other.

X.4.2 State-independent multi utility

Does monotonicity deliver any form of state independence at all in the context of uncertainty averse preferences with imprecise tastes? The main result of this section provides an answer.

To formulate it, recall representation (X.4) (Section X.3.3), which is equivalent to the basic multi evaluation representation (X.1) but involves a function from evaluations to sets of pairs, satisfying certain properties. We shall continue to use such functions (albeit on probabilities rather than evaluations), with properties as defined in Section X.3.3, [19] and two

18. For reasons suggested below, we conjecture that this incompatibility extends to weak ordered continuous preferences exhibiting a larger range of (non-neutral) uncertainty attitudes.

19. To be precise, all properties are translated without change into the current context ($\mathcal{N}^1$ is replaced by Δ, representation (X.4) by (X.11) and so on), except for non-triviality, where the requirement that the element is non-constant is removed.

new properties. We shall say that a function $\alpha : \Delta \to 2^{\Re_{>0}\times\Re}$ is *grounded* if there exists $p \in \Delta$ such that $(1,0) \in \alpha(p)$. It is *calibrated* with respect to a set of utilities $\mathcal{U}^{ut} \subseteq \Re^X$ if, for all $p \in \Delta$, $(a,b) \in \alpha(p)$, and $x \in \mathcal{U}^{ut}(\Delta(X))$,[20] $ax + b \geqslant x$. (Where it is obvious from the context, the set of utilities $\mathcal{U}^{ut}$ will not be mentioned.) Equipped with these notions, we now state the main theorem of this section.

Theorem X.4. *Let $\preceq$ be a preference relation on $\mathcal{A}$. The following are equivalent:*

(i) *$\preceq$ satisfies X-A1–X-A6 and X-A9,*

(ii) *there exist a non-trivial minimal tight closed convex set of utility functions $\mathcal{U}^{ut} \subseteq \Re^X$ and a non-trivial, tight, grounded, calibrated, upper hemicontinuous, convex function $\alpha : \Delta \to 2^{\Re_{\geqslant 0}\times\Re}$ such that $\preceq$ is represented by:*

$$V(f) = \min_{p\in\Delta,\ (a,b)\in\alpha(p)} \left(a \sum_{s\in S} p(s) \min_{u\in\mathcal{U}^{ut}} u(f(s)) + b \right) \tag{X.11}$$

Moreover, for any minimal tight closed convex $\mathcal{U}^{ut\prime} \subseteq \Re^X$ and tight, grounded, calibrated upper hemicontinuous, convex function $\alpha' : \Delta \to 2^{\Re_{\geqslant 0}\times\Re}$ representing $\preceq$, there exists $a \in \Re_{\geqslant 0}$ and $b \in \Re$ such that $\mathcal{U}^{ut\prime} = a\mathcal{U}^{ut} + b$ and $\alpha' = \alpha$.

Monotonicity does imply a separation of tastes for outcomes and beliefs: the former are represented by a set of (state-independent) utilities and the latter are incorporated into a function that only operates at the level of states (and probabilities measures over them). Moreover, this separation does involve a form of state independence: the same set of utility functions is used to assess the consequence of every act in every state. It is this set – this multi utility – that is independent of the state. We shall thus refer to (X.11) as a *state-independent multi utility (multi prior) representation.*

In order to elucidate the relationship with the notion of multi state-independent utility considered in Section X.4.1, it is useful to rewrite representation (X.11) in the same form as (X.1). To this end, for any $\mathcal{U}^{ut} \subseteq \Re^X$, let $(\mathcal{U}^{ut})^S = \{u : S \times X \to \Re \mid \forall s \in S,\ u(s, \bullet) \in \mathcal{U}\}$. $(\mathcal{U}^{ut})^S$ is the set of state-dependent utilities whose restriction to any state coincides with a utility in $\mathcal{U}^{ut}$.[21] Let $\langle(\mathcal{U}^{ut})^S\rangle = \{a\mathcal{U}^S + b \mid \forall a \in \Re_{>0}, b \in \Re\}$.

20. We continue to use the notation introduced in Section X.2: $\mathcal{U}^{ut}(\Delta(X)) = \{\min_{u\in\mathcal{U}^{ut}} u(c) \mid c \in \Delta(X)\}$.

21. Note that, unless $\mathcal{U}^{ut}$ is a singleton, $(\mathcal{U}^{ut})^S$ will contain at least one state-dependent utility.

A set $\mathcal{U} \subseteq \Re^{S\times X}$ is *rectangular* if i. there exist $\mathcal{U}^{ut} \subseteq \Re^X$, $\mathcal{C} \subseteq \Delta$ such that $\mathcal{U} \subseteq \left\langle (\mathcal{U}^{ut})^S \right\rangle \times \mathcal{C}$ and ii. for all $p \in \mathcal{C}$ and $U \in (\mathcal{U}^{ut})^S$, if $(p, au + b) \in \mathcal{U}'$ for some $a \in \Re_{>0}$, $b \in \Re$, then $(p, au' + b) \in \mathcal{U}$ for every $u' \in (\mathcal{U}^{ut})^S$. A rectangular set is, in an appropriately generalised sense, a product of a set of probabilities and the set of state-dependent utilities that are 'generated' from a set of utilities. In particular, rectangular sets consist of pairs of probability measures and state-dependent utility functions, with an added constraint that any probability that appears somewhere in the set is matched with every state-dependent utility that appears somewhere.

It is straightforward to show that any preference relation can be represented according to (X.11) if and only if it can be represented according to (X.1) by a minimal tight rectangular convex closed set $\mathcal{U}^{mon} \subseteq \Re^{S\times X}$. $\mathcal{U}^{mon}$ is a set of probability-utility pairs, but the utilities will not all be state independent in general. State independence at the level of sets of utilities and probabilities – as embodied in state-independent multi utility representations – thus yields state dependence at the level of probability-utility pairs, except in very special cases. Hence the incompatibility reported in Proposition X.2.

The contrast may be illustrated on a continuation of the example given in Section X.4.1.3. Using the same setup and notation, consider a state-independent multi utility representation (X.11) with as set of utilities $\mathcal{U}^{ut}$ the convex closure of $\{u_1, u_2\}$ and with sole probability p.[22] It is straightforward to check that $c \sim d$ and $c_s d_{s^c} \sim c$, as monotonicity demands. This is because, in the assessment of $c_s d_{s^c}$, the value assigned to the consequence c obtained in s is 0 as given by u_1, but the value assigned to the consequence d obtained in t is also 0, following u_2. So a state-dependent utility, which agrees with u_1 in state s and with u_2 in state t, is essentially involved in the assessment of this act under the state-independent multi utility representation. This is the sense in which this representation is state dependent.

On the other hand, the multi state-independent utility representation (X.2) with $\mathcal{U}^{si}$ (see Section X.4.1.3) assigns a value higher than 0 to $c_s d_{s^c}$ because it uses the same state-independent utility to assess c in s and d in t (and when this utility gives a low value to one, it gives a high value to the other). This means in particular that whilst the utility u_1, giving value 0 to c, is used to assign a value to the constant act c, it is not used to assign a value to the consequence c at state s in the context of the act $c_s d_{s^c}$. The *set* of utilities involved

22. More precisely, the α is defined as follows: $\alpha(p) = \{(1, 0)\}$, and $\alpha(q) = \varnothing$ for $q \neq p$.

in the assessment of consequences under the multi state-independent utility representation may depend on the state and the act. In this sense, the representation is state dependent.

Remark X.1. Like representation (X.1), representation (X.11) generalises many of the main theories of uncertainty averse preferences proposed in the literature (see references in Section X.3.2). It does not assume a single utility function – characterised axiomatically by the restriction of independence to risky prospects [23] – whereas, as already mentioned, the vast majority of ambiguity theories do. Moreover, the function α generalises similar 'ambiguity indices' and 'confidence functions' in the variational and confidence models respectively (Maccheroni et al., 2006; Chateauneuf and Faro, 2009). It is relatively straightforward to show that, in the special case of a singleton $\mathcal{U}^{ut}$:

— preferences represented by (X.11) are variational (Maccheroni et al., 2006) whenever $a = 1$ for all $p \in \Delta$ and $(a, b) \in \alpha(p)$;

— preferences represented by (X.11) are confidence preferences (Chateauneuf and Faro, 2009) whenever $\min \mathcal{U}^{ut}(\Delta(X)) = 0$ and $b = 0$ for all $p \in \Delta$ and $(a, b) \in \alpha(p)$; [24]

— preferences represented by (X.11) are maxmin EU preferences (Gilboa and Schmeidler, 1989) when there exists a closed convex $\mathcal{C} \subseteq \Delta$ with $\alpha(p) = \{(1, 0)\}$ for all $p \in \mathcal{C}$ and $\alpha(p) = \varnothing$ for all $p \notin \mathcal{C}$.

Naturally, using the general form of (X.11) under any of these specifications yields multi utility versions of variational, confidence and maxmin EU preferences respectively. Axiomisation of these special cases involves relatively straightforward translations of the results in the aforementioned papers, combined with the insights in the proof of Theorem X.4, and are omitted.

Finally, when there is a single utility function, representation (X.11) becomes a special case of the 'uncertainty averse preferences' characterised by Cerreia-Vioglio et al. (2011b), which are represented by a functional of the following form:

$$V(f) = \inf_{p \in \Delta} G(\sum_{s \in S} p(s)u(f(s)), p) \tag{X.12}$$

23. This axiom is dubbed Risk Independence by Cerreia-Vioglio et al. (2011b,a).

24. If one removes the condition that $\min \mathcal{U}^{ut}(\Delta(X)) = 0$, one obtains the more general class of homothetic preferences, as Cerreia-Vioglio et al. (2011b) call them. They do not involve Chateauneuf and Faro's assumption of a minimal element.

It is straightforward to see that representation (X.11) with singleton $\mathcal{U}^{ut}$ corresponds to the special case where $G(t,p) = \min_{(a,b)\in\alpha(p)} at + b$. By duality, it is thus evident that representation (X.11) with singleton $\mathcal{U}^{ut}$ corresponds to the subclass of Cerreia-Vioglio et al's uncertainty averse preferences where G is concave in the first coordinate. As noted in Section X.3.2, all major theories of uncertainty averse preferences, except for the general case of Cerreia-Vioglio et al. (2011b), belong to this class.

X.5 Related literature

As suggested previously, the results presented here tie into several different literatures. The relationship to some of these literatures has been discussed already: for example, the ambiguity literature in Sections X.3.2 and X.4.2, the state-dependent utility literature in the Introduction and Section X.3.2. The literature on multi utility multi prior representations under incomplete preferences focusses on two of the three issues brought up in the Introduction (imprecise tastes and state dependence), and is discussed at various points in the paper. Beyond Galaabaatar and Karni (2013), whose relation to the work here has been explained in detail, Ok et al. (2012), who characterise a multi state-independent utility single prior representation, is also relevant.

Another related literature involves multi utility-style representations of preferences under risk. Maccheroni (2002) proposes a maxmin multi utility representation of preferences under risk, of which representation (X.7) is a generalisation. Moreover, Corollary X.1 is comparable to the result for convex preferences under risk obtained by Cerreia-Vioglio (2009). Using a conditions reminiscent of our X-A1–X-A4, he obtains a representation by the infimum of a function of (suitably normalised) utility functions and the expected utility calculated with these functions. Rewriting (X.7) in a way similar to (X.4), it becomes clear that it is a special case of representation (2) in Cerreia-Vioglio (2009), with a specific function of utilities and expected utility values.

A final piece of related literature, at least technically, concerns the existence and uniqueness of representations of preference relations by concave functions. The proof of Theorem X.1 draws on a characterisation result provided by Kannai (1977). This contribution and that of Debreu (1976) investigate uniqueness properties that play a role in our results.

X.6 Conclusion

There are apparently many important real-life decisions in which decision makers have difficulty forming precise probabilities, but also furnishing precise utilities. This paper proposes a theory of uncertainty aversion decision making that incorporates both of these imprecisions. Technical axioms aside, the basic model, which can also accommodate state dependence of utility, results from dropping the independence and monotonicity axioms from standard axiomatisations of subjective expected utility.

The central finding of the paper is the subtlety of the meaning of state independence of utility in this context. If tastes are imprecise, and beliefs sufficiently non-trivial, there turn out to be two *incompatible* notions of state independence of utility. One, which is popular in the burgeoning literature on incomplete preferences, involves sets of probability-utility pairs, and imposes state independence on each pair. The other, which corresponds to a standard state-independence axiom, requires state independence at the level of the sets of utilities used in the assessment of the various acts at the different states. The paper provides axiomatic characterisations of these different notions, thus identifying the behavioural differences between them.

X.A Proofs

Throughout the Appendix, $\leqslant$ on $\Re^n$ is the standard order, given by $a \leqslant b$ iff $a_i \leqslant b_i$ for all $1 \leqslant i \leqslant n$, for all $a, b \in \Re^n$. $\cdot$ is the standard scalar product of vectors.

X.A.1 Proofs of results in Section X.3

Proof of Theorem X.1. Consider firstly the (i) implies (ii) direction. X-A1, X-A2 X-A3, X-A4, along with (Cerreia-Vioglio et al., 2011b, Lemma 56) imply that $\preceq$ is a non-degenerate continuous convex complete preference order. X-A5 implies that, for each f that is not maximal under $\preceq$, there is no non-empty open subset of $\{g \in \mathcal{A} \mid g \sim f\}$.

For every $f_1, f_2, f_3 \in \mathcal{A}$ such that $f_1 \prec f_2 \prec f_3$, let

(X.13)
$$\alpha(f_1, f_2, f_3) = \sup \left\{ \frac{|g_{1\beta}g_3 - g_1|}{|g_3 - g_{1\beta}g_3|} \mid g_i \in \mathcal{A}, \beta \in [0,1],\ g_1 \sim f_1, g_3 \sim f_3, g_{1\beta}g_3 \sim f_2 \right\}$$

Note that $\mathcal{A}$ is a convex closed bounded subset of $\Re^{S\times X}$. For any $\hat{x} \in X$, $\mathcal{A}$ is isomorphic to a closed bounded subset of $\Re^{S\times X\backslash\{\hat{x}\}}$ which has non-empty interior. [25] By (Kannai, 1977, Theorem 2.4 and Remark 2.7), the following condition is necessary and sufficient for the existence of a concave continuous functional representing $\preceq$.

Condition 1 There exists a dense sequence of finite sequences of elements of $\mathcal{A}$, $\{f_{i,n}\}$, $0 \leqslant i \leqslant n$, where $f_{0,n} \prec f_{1,n} \prec \cdots \prec f_{n,n}$, $f_{0,n}$ is minimal with respect to $\preceq$ and $f_{n,n}$ is maximal, such that for every $f \in \mathcal{A}$ such that, when n is large enough, there exists $j = j(f,n)$ with $f_{j,n} = f$: [26]

$$\lim_{n\to\infty} \left[\sum_{k=1}^{n} \prod_{i=k}^{n-1} \alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) \right]^{-1} \sum_{k=1}^{j} \prod_{i=k}^{n-1} \alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) < 1 \quad \text{(X.14)}$$

25. Explicitly: the image of $f \in \mathcal{A}$ is $\hat{f} \in \Re^{S\times X\backslash\{\hat{x}\}}$, where $\hat{f}(s,x) = f(s,x)$ for all $s \in S$, $x \in X\backslash\{\hat{x}\}$.

26. Here we adopt the same conventions about sums and products as Kannai (1977): the summation over an empty set of indices is equal to 0, and the product is equal to 1. We also use his definition of a dense sequence of sequences: for every $g, h \in \mathcal{A}$ such that $g \prec h$, there exists N such that for every $n \geqslant N$, there exists i with $g \prec f_{i,n} \prec h$.

Claim X.1. *Condition 1 holds.*

Proof. Let g be a minimal element with respect to $\preceq$ and h be a maximal element (such elements exists because $\mathcal{A}$ is compact as a subset of $\Re^{S\times X}$ and $\preceq$ satisfies X-A3). For each $l \in \mathcal{A}$, $g \prec l \prec h$ with $\alpha(g,l,h) < 1$ (note that by definition $\alpha(g,l,h) \neq 0$), define the following sequence by induction.

— $f^l_0 = g$, $f^l_1 = l$

— suppose f^l_i has been defined. If $\alpha(f^l_{i-1}, f^l_i, h) \leqslant 1$, then, by the continuity of $\preceq$, there exists $h' \succ f^l_i$ such that $\alpha(f^l_{i-1}, f^l_i, h') = 1$; in this case, let $f^l_{i+1} = h'$. If $\alpha(f^l_{i-1}, f^l_i, h) > 1$, then stop the construction; f^l_i is the last element of the sequence.

By construction, $\alpha(f^l_{i-1}, f^l_i, f^l_{i+1}) = 1$ for all i in this sequence. For a given l, let the length of the constructed sequence be n_l, and let $h_l = f^l_{n_l}$. It is straightforward to show that for $f_1 \prec f_2 \prec f_2' \prec f_3$, $\alpha(f_1, f_2, f_3) \leqslant \alpha(f_1, f_2', f_3)$ (it suffices to consider the triple where the supremum in (X.13) is reached; such a triple exists because $\mathcal{A}$ is compact). Moreover, for acts f_1, f_2 and sequences $(f_2)_j$ and $(f_3)_j$ with $(f_2)_j \to f_2$, if $f_1 \prec (f_2)_j, f_2 \prec (f_3)_j$ and $\alpha(f_1, (f_2)_j, (f_3)_j) = 1$ for all j, then by the continuity of $\preceq$ and the fact that $\mathcal{A}$ is compact, there exists f_3 and $(f_3)'_j \sim (f_3)_j$ for all j such that $(f_3)'_j \to f_3$ and $\alpha(f_1, f_2, f_3) = 1$. Hence it follows from the construction that, for any l_j, l_k, if $l_j \prec l_k$ then $n_{l_j} \geqslant n_{l_k}$ and if $l_j \prec l_k$ and $n_{l_j} = n_{l_k} = n$, then $f^{l_j}_n \prec f^{l_k}_n$. Moreover, if $l_j \to l$, then $f^{l_j}_i \to f^l_i$ for each i where this element exists.

Take any $n \geqslant 2$, and take any l such that $n_{l'} \geqslant n$ for all $l' \prec l$ and $n_{l'} < n$ for all $l' \succ l$; it is clear from the construction that such l exists. By the continuity property mentioned above, since l is a the limit of a sequence l_j with $l_j \prec l$ and $n_{l_j} \geqslant n$ for all j, $n_l \geqslant n$. By the continuity property again, for any sequence l_j such that $l_j \succ l$ and $n_{l_j} \leqslant n-1$ for all j, and $l_j \to l$, $\alpha(f^{l_j}_{n-2}, f^{l_j}_{n-1}, h) \to \alpha(f^l_{n-2}, f^l_{n-1}, h)$. So if $\alpha(f^l_{n-2}, f^l_{n-1}, h) < 1$, then there exists $l_j \succ l$ such that $\alpha(f^{l_j}_{n-2}, f^{l_j}_{n-1}, h) < 1$, and hence, by construction, there exists an element $f^{l_j}_n$, contradicting the assumption that $n_{l'} < n$ for all $l' \succ l$. Hence $\alpha(f^l_{n-2}, f^l_{n-1}, h) = 1$, and $h \sim f^l_n$. Hence, for any $n \geqslant 2$, there exists a sequence $f_{0,n}, \dots f_{n,n}$ of length n with $f_{0,n} = g$, $f_{n,n} = h$ and $\alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) = 1$ for all $1 \leqslant i \leqslant n-1$. Take a sequence of such sequences, $\{f_{i,n}\}$.

Now we show that this sequence is dense in $\preceq$. Without loss of generality, we may assume that $g \prec h_\alpha g \prec h$ for all $\alpha \in (0,1)$ (if not, replace g with the maximal $h_\beta g$ such that $h_\beta g \preceq g$, and similarly for h; such acts exist by the continuity of $\preceq$). For each sequence $(f_{i,n})_{0\leqslant i\leqslant n}$ in the sequence of sequences, and for each $0 \leqslant i \leqslant n$, let $\alpha_{i,n} = \min\{\beta \in [0,1] \mid h_\beta g \succeq f_{i,n}\}$ (the minimum exists by the continuity of $\preceq$). $(\alpha_{i,n})_{0\leqslant i\leqslant n}$ is a strictly increasing sequence with $\alpha_{0,n} = 0$ and $\alpha_{1n} = 1$. We claim that, for each $\delta > 0$, there exists $N > 0$ such that for all $n \geqslant N$, $\alpha_{i+1,n} - \alpha_{i,n} \leqslant \delta$ for all $0 \leqslant i \leqslant n-1$; it follows straightforwardly that the sequence is dense. (For any $f, f' \in \mathcal{A}$ such that $f \prec f'$, take $\alpha' = \min\{\beta \in [0,1] \mid h_\beta g \succeq f'\}$ and $\alpha = \max\{\beta \in [0,1] \mid \beta \leqslant \alpha',\ h_\beta g \preceq f\}$: it follows from the claim there exists $N > 0$ such that for all $n \geqslant N$, $\alpha_{i+1,n} - \alpha_{i,n} < \alpha' - \alpha$ for all $0 \leqslant i \leqslant n-1$, so there exists $\alpha < \alpha_{i,n} < \alpha'$ and thus $f \prec f_{i,n} \prec f'$.) To establish the claim, note firstly that, for all $1 \leqslant i \leqslant n-1$ and all n, since $\alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) = 1$, $\alpha_{i,n} - \alpha_{i-1,n} \leqslant \alpha_{i+1,n} - \alpha_{i,n}$. So it suffices to show that, for every $\delta > 0$ there exists $N > 0$ such that, for all $n \geqslant N$, $\alpha_{n,n} - \alpha_{n-1,n} \leqslant \delta$. Suppose that this is not the case, and that $\alpha_{n,n} - \alpha_{n-1,n} \geqslant \epsilon > 0$ for all sufficiently large n. Consider $h_{1-\epsilon}g$; since $\epsilon > 0$, $g \prec h_{1-\epsilon}g \prec h$. By the assumption, for every sequence $(f_{i,n})_{0\leqslant i\leqslant n}$, $\frac{\max_{f_{j,n}\preceq h_{1-\epsilon}g} j}{n} = \frac{n-1}{n}$. So $\frac{\max_{f_{j,n}\preceq h_{1-\epsilon}g} j}{n} \to 1$, contradicting X-A6. Hence for every $\delta > 0$ there exists $N > 0$ such that, for all $n \geqslant N$, $\alpha_{n,n} - \alpha_{n-1,n} \leqslant \delta$, and the sequence of sequences $\{f_{i,n}\}$ is dense as required.

For $f_1 \prec f_2 \prec f_3$, since $\mathcal{A}$ is compact, the supremum in (X.13) is obtained; let $f_1' \sim f_1$, $f_3' \sim f_3$ and $(f_1')_\beta f_3' \sim f_3$ be a triple where the supremum is attained. Since $\frac{|(f_1')_\beta f_3' - f_1'|}{|f_3' - (f_1')_\beta f_3'|} = \frac{\beta |f_3' - f_1'|}{(1-\beta)|f_3' - f_1'|} = \frac{\beta}{(1-\beta)}$, if $\alpha(f_1, f_2, f_3) = 1$, then $\beta = \frac{1}{2}$. It follows that, if $\alpha(f_1, f_2, f_3) = 1$, then f_2 is a conservative potential fifty-fifty mixture of f_1 and f_3. Hence, for any sequence in the sequence of sequences $\{f_{i,n}\}$ defined above, $f_{i,n}$ is a conservative potential fifty-fifty mixture of $f_{i-1,n}$ and $f_{i+1,n}$ for all $0 \leqslant i \leqslant n$. For each such sequence and each $f = f_{j,n}$, $\sum_{k=1}^{j} \prod_{i=k}^{n-1} \alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) = j = \max_{f_{i,n}\preceq f} i$ and $\sum_{k=1}^{n} \prod_{i=k}^{n-1} \alpha\left(f_{i-1,n}, f_{i,n}, f_{i+1,n}\right) = n$. So, if f is such that there exists j for each n with $f = f_{j,n}$, the expression in the limit in (X.14) is $\frac{\max_{f_{i,n}\preceq f} i}{n}$. X-A6 implies that this expression is bounded away from 1, so the limit in (X.14) is less than 1, as required.

□

By Kannai (1977, Theorem 2.4 and Remark 2.7), there exists a concave continuous non-constant functional $V : \mathcal{A} \to \Re$ representing $\preceq$. In fact, as noted by (Kannai, 1977, p11), there exists a minimally concave functional (see also Debreu (1976)) in following sense: for any concave functional V' representing $\preceq$ such that $V'(h) = V(h)$ and $V'(g) = V(g)$ where g and h are minimal and maximal elements with respect to $\preceq$ respectively,

$V'(f) \geqslant V(f)$ for all $f \in \mathcal{A}$. We henceforth assume that V is minimal in this sense.

Let $\hat{x}$ be a generic element of X. An evaluation $U \in \Re^{S \times X}$ is canonical if $U(s, \hat{x}) = 0$ for all $s \in S$. Let $\mathcal{CE}$ be the set of canonical evaluations. Note that for each evaluation, there exists a cardinally equivalent canonical evaluation.

Since $\mathcal{A}$ is a closed bounded convex subset of a finite-dimensional space, and V is concave and finite, for each $f \in ri(\mathcal{A})$,[27] there exists a canonical evaluation supporting V at f, ie. $U \in \mathcal{CE}$ such that $U \cdot g \geqslant V(g)$ for all $g \in \mathcal{A}$ and $U \cdot f = V(f)$. Let $\mathcal{U}$ be convex closure of the set of all canonical evaluations supporting V at some $f \in ri(\mathcal{A})$. By construction, $V(f) = \min_{U \in \mathcal{U}} U \cdot f$ for all $f \in ri(\mathcal{A})$; by the continuity of V, this establishes representation (X.1). By construction $\mathcal{U}$ is closed and convex. Since there is a dense subset of $ri(\mathcal{A})$ on which V is differentiable – and hence at which it has unique supergradients – and the supergradients at all other points are contained in the convex closure of the supergradients at points where V is differentiable (Rockafellar, 1970, Theorems 25.5 and 25.6), $\mathcal{U}$ is tight. That $\mathcal{U}$ is minimal follows from the minimality of V and the following lemma.

Lemma X.A.1. *Let V be a concave continuous function representing $\preceq$, and let $\mathcal{U} = cl(conv\{U \in \mathcal{CE} \mid \forall g \in \mathcal{A}\ U(g) \geqslant V(g),\ \exists f \in \mathcal{A},\ U(f) = V(f)\}))$.[28] Then $\mathcal{U}$ is a minimal (as a tight closed convex set representing $\preceq$) if and only if V is minimal (as a concave representation).*

Proof. Suppose that V is a minimally concave. Consider any tight closed convex $\mathcal{U}'$ representing $\preceq$ according to (X.1) such that $\mathcal{U}'(\mathcal{A}) = \mathcal{U}(\mathcal{A})$, and let $V'(f) = \min_{U \in \mathcal{U}'} U \cdot f$ for all $f \in \mathcal{A}$. Let g, h be minimal and maximal elements of $\mathcal{A}$ with respect to $\preceq$ (such elements exist by X-A3 and the fact that $\mathcal{A}$ is compact); it follows from the assumptions that $V'(g) = V(g)$ and $V'(h) = V(h)$. So, by the minimality of V, $V'(f) \geqslant V(f)$ for all $f \in \mathcal{A}$. Consider any $f \in \mathcal{A}$ and $U' \in \mathcal{U}'$. By the representation, $U' \cdot f \geqslant V'(f)$. However, by the representation, there exists $U \in \mathcal{U}$ such that $U \cdot f = V(f) \leqslant V'(f)$, so $U' \cdot f \geqslant U \cdot f$. Since this holds for every $f \in \mathcal{A}$ and $U' \in \mathcal{U}'$, $\mathcal{U}$ is minimal.

Now suppose that $\mathcal{U}$ is minimal. Let g, h be minimal and maximal elements of $\mathcal{A}$ with respect to $\mathcal{A}$. Let V' be any concave continuous functional representing $\preceq$ with $V'(g) = V(g)$ and $V'(h) = V(h)$, and let $\mathcal{U}' = cl(conv(\{U \in \mathcal{CE} \mid \forall g \in \mathcal{A}\ U(g) \geqslant V(g),\ \exists f \in$

27. For a set X, $ri(X)$ is the relative interior of X.
28. For a set X, $cl(X)$ is its closure and $conv(X)$ its convex hull.

$\mathcal{A}$, $U(f) = V'(f)\}$)). Noting that $\mathcal{U}'(\mathcal{A}) = V'(\mathcal{A}) = V(\mathcal{A}) = \mathcal{U}(\mathcal{A})$, and that $\mathcal{U}'$ is a tight closed convex set representing $\preceq$, minimality of $\mathcal{U}$ implies that for each $U' \in \mathcal{U}'$ and $f \in \mathcal{A}$, there exists $U \in \mathcal{U}$ with $U' \cdot f \geqslant U \cdot f$. Hence $\min_{U' \in \mathcal{U}'} U' \cdot f \geqslant \min_{U \in \mathcal{U}} U \cdot f$ for all $f \in \mathcal{A}$, so $V'(f) \geqslant V(f)$ for all $f \in \mathcal{A}$. Hence V is minimally concave, as required.

□

Necessity The implication (ii) to (i) is standard for all axioms except X-A5 and X-A6. Suppose that $\preceq$ is represented by $\mathcal{U}$ according to (X.1) and let $V(f) = \min_{U \in \mathcal{U}} U \cdot f$ for all $f \in \mathcal{A}$.

For X-A5, consider any $f, g \in \mathcal{A}$ with $g \succ f$. By the representation, for all $U \in \mathcal{U}$, $U \cdot f \geqslant V(f)$ and $U \cdot g \geqslant V(g)$, so $U \cdot f_\alpha g \geqslant \alpha V(f) + (1-\alpha)V(g)$. Since $g \succ f$, it follows that $g_\alpha f \succ f$ for all $\alpha \in (0, 1]$, as required.

As concerns X-A6, note that for any finite sequence $\{f_i\}$ of the sort described in the axiom, $\frac{\max_{f_i \preceq f} i}{n} < 1$, so it needs to be shown that there do not exist any sequence of such sequences $\{f_{i,n}\}$, each of length n, with $\frac{\max_{f_{i,n} \preceq f} i}{n} \to 1$ as $n \to \infty$. By Kannai (1977, Theorem 2.4 and Remark 2.7), if $\preceq$ is representable by a concave functional, then for all $g \prec f \prec h$ and for all dense sequences of sequences $\{f_{i,n}\}$ from g to h satisfying the properties described in Condition 1, (X.14) holds. By the argument in the proof of Claim X.1, for any $g \prec h$, any sequence of sequences $\{f_{i,n}\}$ with $f_{i,n}$ being a conservative potential fifty-fifty mixture of $f_{i-1,n}$ and $f_{i+1,n}$ for all consecutive triples of elements is a dense sequence of sequences from g to h; hence (X.14) holds. Since, as noted in the proof above, the expression in (X.14) simplifies to $\frac{\max_{f_{i,n} \preceq f} i}{n}$ in the case of sequences involving conservative potential fifty-fifty mixes, it follows that $\frac{\max_{f_{i,n} \preceq f} i}{n} \nrightarrow 1$. So X-A6 holds.

Uniqueness Suppose that $\mathcal{U} \neq \mathcal{U}'$ are non-trivial minimal tight closed convex sets representing $\preceq$. Let g, h be minimal and maximal elements of $\preceq$ respectively. Let $V(f) = \min_{U \in \mathcal{U}} U \cdot f$ and $V'(f) = \min_{U' \in \mathcal{U}'} U' \cdot f$ for all $f \in \mathcal{A}$. Let $a \in \Re_{>0}$ and $b \in \Re$ be such that $aV'(g) + b = V(g)$ and $aV'(h) + b = V(h)$ (it is straightforward to show that such a and b exist), and let $\mathcal{U}'' = a\mathcal{U}' + b$ and $V'' = aV' + b$. If $V = V''$, then since $\mathcal{U}$ is the convex closure of the set of supergradients of V, for each member of $\mathcal{U}$, there exists a cardinally equivalent member of $\mathcal{U}''$. Since both sets are tight, it follows that $\mathcal{U}$ and $\mathcal{U}''$ are cardinally equivalent. It thus suffices to show that $V'' = V$. By the minimality of $\mathcal{U}$ and Lemma X.A.1, $V''(f) \geqslant V(f)$ for all $f \in \mathcal{A}$. By the minimality of $\mathcal{U}'$ and Lemma

X.A.1, $V(f) \geqslant V''(f)$ for all $f \in \mathcal{A}$. Hence $V'' = V$, and $\mathcal{U} = a\mathcal{U}' + b$ as required. □

Proof of Proposition X.1. Let $\mathcal{U}$ represent $\preceq$, and let $V(f) = \min_{U \in \mathcal{U}} U \cdot f$ for all $f \in \mathcal{A}$. Let $\preceq_{\Delta(X)}$ be the restriction of $\preceq$ to $\Delta(X)$. For each $c \in \Delta(X)$, there exists a $U \in \mathcal{U}$ such that U supports V at c. Evidently, $\sum_{s \in S} U(s, \bullet)$ supports $V|_{\Delta(X)}$ at c. Let $\mathcal{U}_{\Delta(X)} = cl(conv(\{\sum_{s \in S} U(s, \bullet) \,|\exists c \in \Delta(X) \; s.t. \; U \text{ supports } V \text{ at } c\}))$. Note that, since $\mathcal{U}$ is closed and convex, if $u \in \mathcal{U}_{\Delta(X)}$, there exists $U \in \mathcal{U}$ such that $u = \sum_{s \in S} U(s, \bullet)$. $\mathcal{U}_{\Delta(X)}$ is closed and convex; by arguments similar to those used in the proof of Theorem X.1, it is tight. Uniqueness follows by a similar argument to that used in the proof of Theorem X.1. □

X.A.2 Proofs of Results in Section X.4

Proof of Theorem X.2. First consider the (ii) to (i) direction. The necessity of all axioms apart from X-A7 follows from Theorem X.1. As concerns this axiom, let $\mathcal{U}^{si}$ represent $\preceq$ according to (X.2) and let $V(f) = \min_{(u,p) \in \mathcal{U}^{si}} (u.p) \cdot f$ for all $f \in \mathcal{A}$.

First note that if $s \in S$ is locally null at $g \in \mathcal{A}$ and there is a unique $(u, p) \in \mathcal{U}^{si}$ supporting V at g, then since $(u.p) \cdot (g(s) + \delta e)_s g_{s^c} = \delta p(s) u \cdot e + (u.p) \cdot g = (u.p) \cdot g$ for all $e \in \mathcal{ZUR}$ and $\delta > 0$ sufficiently small, either $p(s) = 0$ or u is constant. By the closure properties of supergradients (Rockafellar, 1970, Theorem 25.6), $Null(s)$ contains all and only acts g where there exists a $(u, p) \in \mathcal{U}^{si}$ supporting V at g with $p(s) = 0$ or u constant. In particular, if $f \notin Null(s)$, then by the continuity and closure properties of supergradients (Rockafellar, 1970, Theorem 25.5 and 25.6), any $(u, p) \in \mathcal{U}^{si}$ supporting V at f is such that $p(s) \neq 0$ and u is not constant. So if s is locally nonnull at f then for every $(u, p) \in \mathcal{U}^{si}$ supporting V at f, $p(s) \neq 0$ and u is not constant.

Consider $f \in \mathcal{A}$ and $s, t \in S$ locally nonnull at f. By the previous observation, for every $(u, p) \in \mathcal{U}^{si}$ supporting V at f, $p(s) \neq 0$ and $p(t) \neq 0$. For $e \in \mathcal{ZUR}$ and $\epsilon > 0$ small enough, $f(s)_s^{e,\delta} f_{s^c} \succeq f$ for all $\delta \leqslant \epsilon$ if and only if $(u.p) \cdot (f(s) + \delta e)_s f_{s^c} = \delta p(s) u \cdot e + (u, p) \cdot f \geqslant (u.p) \cdot f$ for all $\delta \leqslant \epsilon$ and every (u, p) supporting f at V. This holds if and only if $u \cdot e \geqslant 0$ for every (u, p) supporting f at V. Similarly, $f(t)_t^{e,\delta} f_{t^c} \succeq f$ for all $\delta \leqslant \epsilon$ for some ϵ small enough, if and only if $u \cdot e \geqslant 0$ for every (u, p) supporting f at V. Hence there exists $\epsilon > 0$ such that $f(s)_s^{e,\delta} f_{s^c} \succeq f$ for all $\delta \leqslant \epsilon$ if and only if $f(t)_t^{e,\delta} f_{t^c} \succeq f$ for all $\delta \leqslant \epsilon$, and X-A7 is satisfied.

Now consider the (i) to (ii) implication. By Theorem X.1, there is a representation of $\preceq$ by a non-trivial minimal tight closed convex $\mathcal{U} \subseteq \Re^{S\times X}$ according to (X.1). We will say that $U \in \Re^{S\times X}$ decomposes into $(u,p) \in \Re^X \times \Delta$ if, for all $g \in \mathcal{A}$, $\sum_{s\in S}\sum_{x\in X} U(s,x)g(s,x) = \sum_{s\in S}\sum_{x\in X} p(s)u(x)g(s,x)$. It suffices to show that every $U \in \mathcal{U}$ decomposes into a state-independent utility-probability pair.

Let $V(f) = \min_{U\in\mathcal{U}} U \cdot f$ for all $f \in \mathcal{A}$. For each $f \in \mathcal{A}$ and $s \in S$, let $\Delta(X)_{f,s} = \{c_s f_{s^c} \mid c \in \Delta(X)\}$. Let $UCS_{f,s} = \{g \in \Delta(X)_{f,s} \mid V(g) \geqslant V(f)\}$, the upper contour set of f in $\Delta(X)_{f,s}$. Finally, for any affine functional $\phi : \Re^X \to \Re$, we denote, with slight abuse of notation, the corresponding functional on $\Delta(X)_{f,s}$ – that is, the functional taking the value $\phi(c)$ on $c_s f_{s^c}$ – by ϕ.

We first establish the following lemma.

Lemma X.A.2. *For every $f \in ri(\mathcal{A})$ and every $s,t \in S$ locally nonnull at f, if $\{g \mid \phi(g) = \Gamma\}$ is a hyperplane supporting $UCS_{f,s}$ in $\Delta(X)_{f,s}$ at f, then there exists Γ' such that $\{g \mid \phi(g) = \Gamma'\}$ supports $UCS_{f,t}$ in $\Delta(X)_{f,t}$ at f.*

Proof. Let $UCS^o_{f,s} = \{d - f(s) \in \Re^X \mid V(d_s f_{s^c}) \geqslant V(f)\}$ be the mapping of the upper contour set of f in $\Delta(X)_{f,s}$ into $\Re^X$ with the vertex shifted to 0 and let $cone(UCS^o_{f,s})$ be the cone generated by it. Note that by X-A4, if $a \in UCS^o_{f,s}$ then $\beta a \in UCS^o_{f,s}$ for every $\beta \in (0,1)$.

Claim X.2. *For every $f \in ri(\mathcal{A})$ and every $s,t \in S$ locally nonnull at f, $cone(UCS^o_{f,s}) = cone(UCS^o_{f,t})$.*

Proof. Suppose, for reductio, that $cone(UCS^o_{f,s}) \neq cone(UCS^o_{f,t})$ and, without loss of generality, that $a \in cone(UCS^o_{f,s}) \backslash cone(UCS^o_{f,t})$. Since $f \in ri(\mathcal{A})$, there exists $\epsilon > 0$ such that, for $s' = s,t$, $(\epsilon a + f(s'))_{s'} f_{s'^c} \in \mathcal{A}$. By the definition of $cone(UCS_{f,s})$, there exists $\epsilon' < \epsilon$ such that $(\delta a + f(s))_s f_{s^c} \succeq f$ for all $\delta \leqslant \epsilon'$. X-A7 implies that $(\delta a + f(t))_t f_{t^c} \succeq f$ for all $\delta \leqslant \epsilon''$, for some $\epsilon'' > 0$, so $a \in cone(UCS^o_{f,t})$, contradicting the assumption. So $cone(UCS^o_{f,s}) = cone(UCS^o_{f,t})$ as required. □

Since the cones coincide, a hyperplane will support $UCS^o_{f,s}$ if and only if it supports $UCS^o_{f,t}$. Since any hyperplane supporting $UCS_{f,s}$ is a translation of one supporting $UCS^o_{f,s}$, this is sufficient to establish the desired lemma. □

We now distinguish two cases.

Firstly, consider $f \in ri(\mathcal{A})$ such that V is differentiable at f, so there is a unique $U \in \mathcal{U}$ supporting V at f. If U is a constant evaluation, then it trivially decomposes; suppose henceforth that this is not the case. Consider $s \in S$. If $U(s, \bullet)$ is constant (ie. $U(s,x) = U(s,x')$ for all $x, x' \in X$), then $U \cdot (f(s)+\delta e)_s g_{s^c} = \delta U(s,e) + U \cdot f = U \cdot f$ for all $e \in \mathcal{ZUR}$ and δ sufficiently small. Since U is the unique support for V at f, this implies that s is locally null at f. If $U(s, \bullet)$ is non-constant, then by similar reasoning, the fact that U is the unique support at f and the continuity of the superdifferential mapping (Rockafellar, 1970, Theorem 24.4), s is locally nonnull at f. Note that, since U is a non-constant evaluation, there exists a cardinally equivalent U' such that $U'(s,x) = 0$ for all $x \in X$ and all s such that $U'(s, \bullet)$ is constant; we henceforth assume without loss of generality that $U(s, \bullet) = 0$ for all $s \in S$ such that $U'(s, \bullet)$ is constant. Now consider $s \in S$ such that $U(s, \bullet)$ is non-constant. Note that since V is differentiable at f, its restriction to $\Delta(X)_{f,s}$ is differentiable at f, so the hyperplane defined by U, $\{g \in \Delta(X)_{f,s} \mid U(g) = V(f)\}$ is the unique support for $UCS_{f,s}$. Note that this hyperplane can be written as $\{c_s f_{s^c} \mid c \in \Delta(X),\ U(s,c) = \Gamma\}$, for some constant Γ. For every other $t \in S$ such that $U(t, \bullet)$ is non-constant, s and t are locally nonnull, so by Lemma X.A.2, $\{c_t f_{t^c} \mid c \in \Delta,\ U(s,c) = \Gamma'\}$ supports $UCS_{f,t}$ at f for some Γ'. Since, by the previous argument, $\{c_t f_{t^c} \mid c \in \Delta,\ U(t,c) = \Gamma''\}$ also supports $UCS_{f,t}$ for some Γ'', and $UCS_{f,t}$ has a unique supporting hyperplane at f, it follows that $U(t,x) = aU(s,x) + b$ for some $a \in \Re_{>0}$ and $b \in \Re$. This implies that, for each $s \in S$ such that $U(s, \bullet)$ non-constant, there exists $a^s \in \Re_{>0}$ and $b^s \in \Re$ such that $U(s,x) = a^s \sum_{s' \in S} U(s',x) + b^s$. Taking $u(x) = \sum_{s' \in S} U(s',x)$ for all $x \in X$ and $p(s) = \frac{a^s}{\sum_{s' \in S} a^{s'}}$ for s such that $U(s, \bullet)$ non-constant and $p(s) = 0$ otherwise, it is straightforward to show that, for all $g \in \mathcal{A}$, $\sum_{s \in S} \sum_{x \in X} U(s,x) g(s,x) = \sum_{s \in S} \sum_{x \in X} p(s) u(x) g(s,x)$, so U decomposes into (u,p).

Now consider any $f \in ri(\mathcal{A})$ that does not fall into the above case. Note that the set of points where V is differentiable is dense in $ri(\mathcal{A})$ (Rockafellar, 1970, Theorem 25.5), and at each of these points, by the previous case, the supporting U is decomposable into a state-independent utility and probability. Hence, there exists a sequence $(f_i)_{i \geqslant 1}$ with $f_i \to f$ such that, for each $i \geqslant 1$ there exist unique U_i supporting V at f_i, for which $U_i(s, \bullet)$ and $U_i(t, \bullet)$ are positive affine transformations of one another for all $s, t \in S$ for which they are not constant. Hence U, the limit of U_i, supports V at f (Rockafellar, 1970, Theorem 24.4) and is such that $U(s,\bullet)$ and $U(t,\bullet)$ are positive affine transformations of one another for all $s,t\in S$ for which they are not constant; so, by the previous argument, U decomposes into a state-independent-probability pair. Moreover, every support (or supergradient) of V at f is a convex combination of such limits of sequences (Rockafellar, 1970, Theorem 25.6); hence every U supporting V at f decomposes into a state-independent utility-probability pair.

Hence every $U\in\mathcal{U}$ supporting V at a point in $ri(\mathcal{A})$ decomposes into a state-independent utility-probabilty pair. By the tightness of $\mathcal{U}$, $\mathcal{U}$ is the convex closure of the set of such U. By the argument above, it follows that every element of $\mathcal{U}$ decomposes into a state-independent utility-probabilty pair, as required. Let $\mathcal{U}^{si}$ be the set of such pairs, for all $U\in\mathcal{U}$. By construction, this set represents $\preceq$ according to (X.2) and is minimal tight closed and convex.

Uniqueness follows immediately from the uniqueness clause in Theorem X.1 and the decomposition into state-independent utility-probability pairs.

□

Proof of Theorem X.3. The (ii) to (i) direction is straightforward; consider the (i) to (ii) implication. By Theorem X.2, there is a representation of $\preceq$ by a minimal tight convex closed $\mathcal{U}^{si}\subseteq\Re^X\times\Delta$ according to (X.2). Let $V(f)=\min_{(u,p)\in\mathcal{U}^{si}}(u.p)\cdot f$, for all $f\in\mathcal{A}$. Let $\mathcal{U}^{ut}=cl\left(conv\{u\in R^X\mid(u,p)\in\mathcal{U}^{si},\ \exists c\in\Delta(X)\ s.t.\ u \text{ supports } V|_{\Delta(X)} \text{ at } c\}\right)$, and $\mathcal{C}=cl\left(conv\{p\in\Delta\mid\forall f\in\mathcal{A}\ f^p\succeq f\ ,\ \exists g\in ri(\mathcal{A})\ s.t.\ g^p\sim g\}\right)$. We will show that $\mathcal{U}^{ut}$, $\mathcal{C}$ represent $\preceq$ according to (X.9).

For any act $f\in\mathcal{A}$, let $c_f\in\Delta(X)$ be the minimal constant act (in the order $\preceq$) of the form f^p for some $p\in\mathcal{C}$. It follows from X-A8 that, for all $f,g\in\mathcal{A}$, $c_f\succeq c_g$ if and only if $f\succeq g$. Since V represents $\preceq$, it follows that $V(f)\geqslant V(g)$ if and only if $V(c_f)\geqslant V(c_g)$. By the definition of c_f and $\mathcal{U}^{ut}$, for every $f,g\in ri(\mathcal{A})$, $V(c_f)\geqslant V(c_g)$ if and only if $\min_{u\in\mathcal{U}^{ut}}\min_{p\in\mathcal{C}}(u.p)\cdot f\geqslant\min_{u\in\mathcal{U}^{ut}}\min_{p\in\mathcal{C}}(u.p)\cdot g$; by the continuity of V, this equivalence continues to hold for all $f,g\in\mathcal{A}$. This yields the desired representation.

By Proposition X.1 and its proof, $\mathcal{U}^{ut}$ is a tight representation of the restriction of $\preceq$ to $\Delta(X)$. In particular, no proper closed convex subset of $\mathcal{U}^{ut}$ represents the restriction of $\preceq$ to $\Delta(X)$; so, to show that $\mathcal{U}^{ut}$ and $\mathcal{C}$ are tight, it suffices to show that there is no proper closed convex subset of $\mathcal{C}$ which, taken with $\mathcal{U}^{ut}$, represents $\preceq$. If $(u,p)\in\mathcal{U}^{ut}\times\Delta$ supports V at $f\in ri(\mathcal{A})$, then $g^p\succeq g$ for all $g\in A$ and $f^p\sim f$. If moreover V is differentiable at f, then (u,p) is the unique support for V at f. Hence, for all other $p' \in \Delta$ such that $g^{p'} \geq g$ for all $g \in \mathcal{A}$, $f^{p'} \succ f$: if not, there would be another supergradient at f, contradicting the assumption that V is differentiable there. Let $\mathcal{D} = \{p \in \Delta \mid \exists f \in ri(\mathcal{A}), V$ differentiable at $f, \exists u \in \mathcal{U}^{ut}\ s.t.\ (u,p)$ supports V at $f\}$; as just noted, $\mathcal{D} \subseteq \mathcal{C}$. If $(u,p) \in \mathcal{U}^{ut} \times \Delta$ supports V at $f \in ri(\mathcal{A})$ where V is not differentiable, $p \in cl\,(conv(\mathcal{D}))$ (Rockafellar, 1970, Theorem 25.6). Since, for any $p \in \Delta$ such that $g^p \geq g$ for all $g \in A$ and $f^p \sim f$ for some $f \in ri(\mathcal{A})$, there exists $u \in \mathcal{U}^{ut}$ such that (u,p) supports V at f, it follows that $\mathcal{C} = cl\,(conv(\mathcal{D}))$. So there is no convex closed proper subset of $\mathcal{C}$ representing $\preceq$ in tandem with $\mathcal{U}^{ut}$ according to (X.9). So the pair $\mathcal{U}^{ut}$ and $\mathcal{C}$ is tight. Minimality follows from the minimality of $\mathcal{U}^{si}$.

As concerns the uniqueness clause, let $\mathcal{U}^{ut\prime}$ and $\mathcal{C}'$ be another pair representing $\preceq$ according to (X.9). Since $\mathcal{U}^{ut\prime}$ and $\mathcal{U}^{ut}$ both represent the restriction of $\preceq$ to $\Delta(X)$, and are minimal tight convex closed sets, it follows from the uniqueness clause in Theorem X.1 (see also Corollary X.1) that $\mathcal{U}^{ut\prime} = a\mathcal{U}^{ut} + b$ for some $a \in \Re_{>0}$, $b \in \Re$. We may thus assume without loss of generality that $\mathcal{U}^{ut\prime} = \mathcal{U}^{ut}$. Note that since, for each $f \in \mathcal{A}$, there exists $p \in \Delta$ such that $f^p \sim f$, $\mathcal{U}^{ut} \times \mathcal{C}'$ and $\mathcal{U}^{ut} \times \mathcal{C}$ generate the same functional $V : \mathcal{A} \to \Re$ representing $\preceq$ according to (X.9). Since both of these sets contain the convex closure of the set of supergradients of V, and $\mathcal{C}'$ and $\mathcal{C}$ are tight, it follows that $\mathcal{C}' = \mathcal{C}$, as required.

□

Proof of Proposition X.2. Let $\mathcal{U}^{si} \subseteq \Re^X \times \Delta$ be a minimal tight closed convex set representing $\preceq$ according to (X.2), and let V be the corresponding functional. Suppose that X-A9 holds. As a point of notation, for every $s \in S$, let $\delta_s \in \Delta$ be the degenerate measure putting all weight on s: $\delta_s(s) = 1$, $\delta_s(t) = 0$ for $t \neq s$. If $p = \delta_s$ for some $s \in S$, then we shall say that p is *degenerate*. Let $\mathcal{ZR} = \{e \in \Re^X \mid \sum_{x \in X} e(x) = 0\}$. Finally, for $c, c' \in \Delta(X)$ with $c \prec c'$, let $(\mathcal{U}^{si})^{[c,c']} \subseteq \mathcal{U}^{si}$ be a minimal tight closed convex representation of $\preceq_{[c,c']}$ according to (X.2); by the uniqueness clause in Theorem X.2 and the properties of minimality (Kannai, 1977), a unique such set exists.

For $c, c' \in \Delta(X)$ such that $c \prec c'$, let $(c,c') = \{d \in \Delta(X) \mid c \prec d \prec c'\}$; we call this the *interval* from c to c'. An interval (c,c') is *single-utilitied* if there exists non-constant $u \in \Re^X$ such that, for every $d \in (c,c')$, every support for $V|_{\Delta(X)}$ at d is a positive affine transformation of u. When writing sets of intervals $\{(\underline{c_1}, \overline{c_1}), \ldots, (\underline{c_n}, \overline{c_n})\}$, we shall henceforth adopt the convention that the intervals are pair-wise disjoint, and that $\overline{c_i} \preceq \underline{c_{i+1}}$ for all i. For any pair of sets of single-utilitied intervals, the appropriately defined union of them contains single-utilied intervals. There is thus a unique set $\mathcal{I}$ of pair-wise disjoint single-utilitied intervals that is maximal under containment: for any other set $\mathcal{I}'$ of single-utilitied intervals, $\bigcup_{(\underline{c_i},\overline{c_i})\in\mathcal{I}'}(\underline{c_i},\overline{c_i}) \subseteq \bigcup_{(\underline{c_i},\overline{c_i})\in\mathcal{I}}(\underline{c_i},\overline{c_i})$. We shall refer to this set henceforth as $\mathcal{I} = \{(\underline{c_1},\overline{c_1}),\ldots,(\underline{c_n},\overline{c_n})\}$.

By X-A9, there exist $\underline{c},\overline{c} \in \Delta(X)$ such that $\underline{c} \preceq f \preceq \overline{c}$ for all $f \in \mathcal{A}$. Fix any such pair of constant acts $\underline{c},\overline{c}$ for reference throughout the proof. Define the following sequence of elements of $\Delta(X)$: if $\underline{c_1} \sim \underline{c}$ and $\overline{c_n} \sim \overline{c}$, then take the sequence $\underline{c_1},\overline{c_1},\underline{c_2},\ldots,\overline{c_{n-1}},\underline{c_n},\overline{c_n}$; if $\underline{c_1} \succ \underline{c}$ and $\overline{c_n} \prec \overline{c}$, then take the sequence $\underline{c} = \overline{c_0},\underline{c_1},\overline{c_1},\underline{c_2},\ldots,\overline{c_{n-1}},\underline{c_n},\overline{c_n},\underline{c_{n+1}} = \overline{c}$; and similarly for the other cases.

The result will be a consequence of the following two claims.

Claim X.3. *If an interval $(\underline{c_i},\overline{c_i})$ is single-utilitied, then $(\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)}$ is a singleton.*

Claim X.4. *For each i, there exists $f \in \mathcal{A}$ such that $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for all $s \in S$. Moreover, for any such f, if V is differentiable at f with support (u,p), then p is degenerate.*

Proof of Claim X.3. Since $(\underline{c_i},\overline{c_i})$ is single-utilitied, there exists $\hat{u} \in \Re^X$ such that for each $u \in (\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)}$, $u = a\hat{u} + b$ for some $a \in \Re_{>0}$, $b \in \Re$. Note that $\hat{u}$ represents the restriction of $\preceq$ to $(\underline{c_i},\overline{c_i})$: for if $\hat{u}\cdot c \geqslant \hat{u}\cdot d$, then for every $u \in (\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)}$, $u\cdot c \geqslant u\cdot d$, so by representation (X.2), $c \succeq d$, and similarly for strict inequalities.

Let $w : \Delta(X) \to \Re$ be given by: $w(c) = \min_{u\in(\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)}} u\cdot c = \min_{a\hat{u}+b\in(\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)}} a\hat{u}\cdot c + b$. w is obviously a concave functional, which represents the restriction of $\preceq$ to $(\underline{c_i},\overline{c_i})$. We now show that w is linear: if this is the case, then all the (a,b) are equal, so since $(\mathcal{U}^{si})^{[\underline{c_i},\overline{c_i}]}_{\Delta(X)} = cl(conv(\{u \mid \exists c \in (\underline{c_i},\overline{c_i})\ s.t.\ (u,p) \in \mathcal{U}^{si}$ supports V at $c\}))$ (by the proofs of Theorem X.1 and Proposition X.1), it is a singleton. Suppose for reductio that w is not linear. Let $u' = A\hat{u} + B$ be such that $u'(\underline{c_i}) = \inf w((\underline{c_i},\overline{c_i}))$ and $u'(\overline{c_i}) = \sup w((\underline{c_i},\overline{c_i}))$. Note that u' represents the restriction of $\preceq$ to $(\underline{c_i},\overline{c_i})$. Moreover, by the (strict) concavity of w, $u'\cdot d \leqslant w(d)$ for all $d \in (\underline{c_i},\overline{c_i})$ with strict inequality for some $d \in (\underline{c_i},\overline{c_i})$, and $u'\cdot d \geqslant w(d)$ for all $d \notin (\underline{c_i},\overline{c_i})$. Let $v' : \Delta(X) \to \Re$ be defined by: $v'(c) = \min_{u\in\mathcal{U}^{si}_{\Delta(X)}\cup\{u'\}} u(c)$. It is a concave function representing the restriction of $\preceq$ to $\Delta(X)$. Moreover, $V|_{\Delta(X)}(c) \geqslant v'(c)$ for all $c \in \Delta(X)$ with strict inequality for some such $c \in (\underline{c_i},\overline{c_i})$. Hence $V|_{\Delta(X)}$ is not a minimal concave representation of the restriction of $\preceq$ to $\Delta(X)$. This contradicts the fact, established by Lemma X.A.6 in the proof of Theorem X.4, that in the presence of X-A9, $V|_{\Delta(X)}$ is minimal. So w is linear and the claim is established.

□

Proof of Claim X.4. To establish the claim, we require the following two Lemmas, which contain the basic insight at the heart of the whole result.

Lemma X.A.3. *Suppose that V is differentiable at $f \in \mathcal{A}$, with support (u, p), where u is non-constant. Then for all $s \in S$, if $p(s) > 0$, then every support of $V|_{\Delta(X)}$ at $f(s)$ is a positive affine transformation of u.*

Lemma X.A.4. *Let $f \in \mathcal{A}$ be such that V is differentiable at f, with support (u, p), where u is non-constant. Suppose that for $s \in S$, $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for some i. If $p(s) > 0$, then $p = \delta_s$.*

Proof of Lemma X.A.3. Let (u, p) be the unique support of V at some $f \in ri(\mathcal{A})$. Take $s \in S$ such that $p(s) > 0$. Since V is differentiable at f, its restriction to $\{c_s f_{s^c} \mid c \in \Delta(X)\}$ is differentiable at f; by representation (X.2), u is the unique support at this point. We show that every support for $V|_{\Delta(X)}$ at $f(s)$ is a positive affine transformation of u. Suppose that this is not the case, and that u', which is not a positive affine transformation of u, supports $V|_{\Delta(X)}$ at $f(s)$. Hence there exists $c', d' \in \Delta(X)$ such that $u(c') > u(d')$ and $u'(c') \leqslant u'(d')$. So, taking $e = c' - d'$ yields $e \in \mathcal{ZR}$ with $u \cdot e > 0$ and $u' \cdot e \leqslant 0$. Since $f \in ri(\mathcal{A})$, for sufficiently small $\delta > 0$, $f(s) + \delta e \in \Delta(X)$. The preceding facts, combined with the fact that u is the unique support to the restriction of V $\{c_s f_{s^c} \mid c \in \Delta(X)\}$ at f, imply that for sufficiently small $\delta > 0$, $(f(s) + \delta e)_s f_{s^c} \succ f$ whereas $f(s) + \delta e \preceq f(s)$, contradicting X-A9. Hence every support of $V|_{\Delta(X)}$ at $f(s)$ is a positive affine transformation of u. This establishes the claim for $f \in ri(\mathcal{A})$. By the continuity of superdifferentials (Rockafellar, 1970, Theorem 24.4), the claim holds for $f \in \mathcal{A}$ where V has a unique support. □

Proof of Lemma X.A.4. Let f be such that V is differentiable at f, with support (u, p), where u is non-constant, and suppose that there exists $s_1 \in S$ such that $\overline{c_i} \prec f(s_1) \prec \underline{c_{i+1}}$. Suppose for reductio that $p(s_1) > 0$ and $p(s_2) > 0$, for some $s_2 \neq s_1$. Since V is differentiable at f with support (u, p), by Lemma X.A.3 every support of $V|_{\Delta(X)}$ at $f(s_1)$ is a positive affine transformation of u, and similarly for $f(s_2)$.

Take any $d \in \Delta(X)$ such that $d \sim f(s_1)$ and let $g = d_{s_1} f_{s_1^c}$. By representation (X.2), there exists $e_1, e_2 \in \mathcal{ZR}$ and $\epsilon > 0$ such that, for all $\delta < \epsilon$, $(f(s_1) + \delta e_1)_{s_1} f_{s_1^c} \nsim f$,

$(f(s_2)+\delta e_2)_{s_2}f_{s_2^c} \not\sim f$, and $(f(s_1)+\delta e_1)_{s_1}(f(s_2)+\delta e_2)_{s_2}f_{\{s_1,s_2\}^c} \sim f$. By the continuity of $\preceq$, for each $\delta < \epsilon$, there exists $e_1^\delta, e_2^\delta \in \mathcal{ZR}$ such that $g(s_1)+e_1^\delta \sim f(s_1)+\delta e_1$ and $g(s_2)+e_2^\delta \sim f(s_2)+\delta e_2$. By X-A9, it follows that $(g(s_1)+e_1^\delta)_{s_1}g_{s_1^c} \not\sim g$, $(g(s_2)+e_2^\delta)_{s_2}g_{s_2^c} \not\sim g$ but $(g(s_1)+e_1^\delta)_{s_1}(g(s_2)+e_2')_{s_2}g_{\{s_1,s_2\}^c} \sim g$ for every δ (and associated e_1^δ). By representation (X.2), it follows that V is supported at g by (u',p') where $p'(s_1) > 0$ and $p'(s_2) > 0$, and that it is not supported by (u'',p'') with $p''(s_1) > 0$ and $p''(s_2) = 0$ or $p''(s_1) = 0$ and $p''(s_2) > 0$.

By the maximality of $\mathcal{I}$, the density of differentiable points of V and the determination of supergradients by supergradients at differentiable points (Rockafellar, 1970, Theorems 25.5 and 25.6), for every subinterval of $(\overline{c_i}, \underline{c_{i+1}})$ containing $f(s_1)$, there exists $d' \in (\overline{c_i}, \underline{c_{i+1}})$ with V differentiable at $g' = d'_{s_1}f_{s_1^c}$ and $V|_{\Delta(X)}$ not supported at d' by a positive affine transformation of u. Taking successively smaller intervals yields a sequence of d_i with these properties, with $d_i \to d$. Since, as shown above, V is supported at $g = d_{s_1}f_{s_1^c}$ by (u',p') where $p'(s_1) > 0$ and $p'(s_2) > 0$ and it is not supported by (u'',p'') with $p''(s_1) > 0$ and $p''(s_2) = 0$ or $p''(s_1) = 0$ and $p''(s_2) > 0$, by the continuity of the superdifferential mapping (Rockafellar, 1970, Theorem 24.4), for i sufficiently large, V is supported at $g_i = (d_i)_{s_1}f_{s_1^c}$ by (u_i,p_i) with u_i non-constant, $p_i(s) \neq 0$ and $p_i(s') \neq 0$. Consider any such i: since $V|_{\Delta(X)}$ is supported at $f(s_2)$ only by positive affine transformations of u, whereas it is not supported at d_i by any positive affine transformation of u, we have a contradiction by Lemma X.A.3. Hence there does not exist $s_2 \in S$ with $p(s_2) > 0$, so $p = \delta_{s_1}$, as required.

□

Now return to Claim X.4. For any i, consider the interval $(\overline{c_i}, \underline{c_{i+1}})$. By the continuity of preferences, the indifference curve in $\Delta(X)$ containing $\overline{c_i}$ is determined by the unique utility function in $(\mathcal{U}^{si})^{[c_i,\overline{c_i}]}_{\Delta(X)}$ (this set is a singleton by Claim X.3), and similarly for $\underline{c_{i+1}}$. So if $\overline{c_i} = \underline{c_{i+1}}$, these two indifference curves are identical, so the utility representing $(\mathcal{U}^{si})^{[c_i,\overline{c_i}]}_{\Delta(X)}$ is a positive affine transformation of that representing $(\mathcal{U}^{si})^{[c_{i+1},\overline{c_{i+1}}]}_{\Delta(X)}$, contradicting the maximality of $\mathcal{I}$. $(\overline{c_i}, \underline{c_{i+1}})$ is thus non-empty, as required.

As concerns the other part of the claim, note that, by X-A9, for any $f \in \mathcal{A}$ such that $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for all $s \in S$, $f \prec \underline{c_{i+1}}$, so it is not supported by (u,p) with u constant. So, for any such f with V is differentiable at f and having support (u,p) there, it follows directly from Lemma X.A.4 that p is degenerate, as required.

$\square$

Claim X.3 establishes that each section $\preceq_{[\underline{c_i},\overline{c_i}]}$ exhibits precise tastes. We now show that each section $\preceq_{[\overline{c_i},\underline{c_{i+1}}]}$ has rudimentary beliefs. Consider $\preceq_{[\overline{c_i},\underline{c_{i+1}}]}$. By Claim X.4, and the fact that supergradients of V are determined by supergradients at differentiable points, for every $f \in \mathcal{A}$ with $\overline{c_i} \preceq f(s) \preceq \underline{c_{i+1}}$ for all $s \in S$, V is supported at f by some (u,p) with degenerate p. We first show that for every $s \in S$ such that there exists $(u,\delta_s) \in \mathcal{U}^{si}$ supporting V at some f with $\overline{c_i} \prec f(s') \prec \underline{c_{i+1}}$ for all $s' \in S$, $\{u \mid (u,\delta_s) \in \mathcal{U}^{si}\}$ represents the restriction of $\preceq$ to $\{d \in (\underline{c_i},\overline{c_i}) \mid d_s\overline{c}_{s^c} \prec \overline{c}\}$.

Consider any such $s \in S$. By representation (X.2), there exists $d \in (\overline{c_i},\underline{c_{i+1}})$ with $d_s\overline{c}_{s^c} \prec \overline{c}$. We show that, for every $d,d' \in (\overline{c_i},\underline{c_{i+1}})$ with either $d_s\overline{c}_{s^c} \prec \overline{c}$ or $d'_s\overline{c}_{s^c} \prec \overline{c}$, if $d' \succ d$ then $d'_s\overline{c}_{s^c} \succ d_s\overline{c}_{s^c}$. Suppose that this is not the case: by X-A9 it follows that there exist $d,d' \in (\overline{c_i},\underline{c_{i+1}})$ with $d_s\overline{c}_{s^c} \prec \overline{c}$ and $d' \succ d$ but $d'_s\overline{c}_{s^c} \sim d_s\overline{c}_{s^c}$. By X-A9, for every $c \in \Delta(X)$ such that $d' \succeq c \succeq d$, $c_s\overline{c}_{s^c} \sim d_s\overline{c}_{s^c}$. For every $g \in \mathcal{A}$, if $g(s) \succeq d$, then $d \preceq g(s)_\alpha d \preceq d'$ for sufficiently small $\alpha \in (0,1)$. Since, by the aforementioned properties and Lemma X.A.4, there exists $(u',\delta_s) \in \mathcal{U}^{si}$ such that $(u'.\delta_s)\cdot g_\alpha(d_s\overline{c}_{s^c}) = (u'.\delta_s)\cdot(g(s)_\alpha d)_s\overline{c}_{s^c} = V(d_s\overline{c}_{s^c})$, it follows that $g_\alpha(d_s\overline{c}_{s^c}) \preceq d_s\overline{c}_{s^c}$ for sufficiently small $\alpha \in (0,1)$. By X-A9, the same holds for $g \in \mathcal{A}$ with $g(s) \preceq d$. So, for every $g \in \mathcal{A}$, for sufficiently small $\alpha \in (0,1)$, $g_\alpha(d_s\overline{c}_{s^c}) \preceq d_s\overline{c}_{s^c}$, contradicting X-A5. Hence the restriction of $\preceq$ to $\{d_s\overline{c}_{s^c} \mid d \in (\overline{c_i},\underline{c_{i+1}}), d_s\overline{c}_{s^c} \prec \overline{c}\}$ coincides with the restriction of $\preceq$ to $\{d \in (\underline{c_i},\overline{c_i}) \mid d_s\overline{c}_{s^c} \prec \overline{c}\}$. Since, by Lemma X.4 and representation (X.2), $\{u \mid (u,\delta_s) \in \mathcal{U}^{si}\}$ represents the former relation, it also represents the latter.

By Lemma X.A.6 in the proof of Theorem X.4, $V|_{\Delta(X)}$ is a minimal concave representation of the restriction of $\preceq$ to $\Delta(X)$. So, for each $s \in S$ such that there exists $(u,\delta_s) \in \mathcal{U}^{si}$ supporting V at some f with $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for all $s \in S$, since $\{u \mid (u,\delta_s) \in \mathcal{U}^{si}\}$ represents the restriction of $\preceq$ to $\{d \in (\underline{c_i},\overline{c_i}) \mid d_s\overline{c}_{s^c} \prec \overline{c}\}$, there exists a concave function $\psi_s : (V|_{\Delta(X)}(\overline{c_i}), V|_{\Delta(X)}(\underline{c_{i+1}})) \to \Re$ such that $\min_{(u,\delta_s)\in\mathcal{U}^{si}} u(c) = \psi_s \circ V|_{\Delta(X)}(c)$ for all $c \in \Delta(X)$. (ψ_s is constant on $V|_{\Delta(X)}(d \in (\underline{c_i},\overline{c_i}) \mid d_s\overline{c}_{s^c} \sim \overline{c}\})$.) Let $\phi_s : V|_{\Delta(X)}(\Delta(X)) \to \Re$ be any concave function extending ψ_S. Hence, for every act f with $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for all $s \in S$, $V(f) = \min_{s\in E}\phi_s(\min_{u\in\mathcal{U}^{si}_{\Delta(X)}} u(f(s)))$, where E is the set of $s \in S$ such that there exists $(u,\delta_s) \in \mathcal{U}^{si}$ supporting V at some f with $\overline{c_i} \prec f(s) \prec \underline{c_{i+1}}$ for all $s \in S$. By the continuity of $\preceq$, this representation continues to hold for f with $\overline{c_i} \preceq f(s) \preceq \underline{c_{i+1}}$ for all $s \in S$; hence $\preceq_{[\overline{c_i},\underline{c_{i+1}}]}$ involves rudimentary beliefs, as required. $\square$

Proof of Theorem X.4. The (ii) to (i) direction is straightforward; consider the (i) to (ii) implication. Theorem X.1 implies that there is a representation of $\preceq$ by a non-trivial minimal tight closed convex $\mathcal{U} \subseteq \Re^{S\times X}$ according to (X.1). Let $V(f) = \min_{U\in\mathcal{U}} U \cdot f$ for all $f \in \mathcal{A}$; as noted in the proof of Theorem X.1 (see Lemma X.A.1), V is a minimally concave representation, in the sense of (Kannai, 1977, p11). Let $\mathcal{U}^{ut} = \mathcal{U}_{\Delta(X)} = \{\sum_{s\in S} U(s, \bullet) \mid \exists c \in \Delta(X)\ s.t.\ U$ supports V at $c\}$. By Proposition X.1 and its proof, $\mathcal{U}^{ut}$ is a tight closed convex set of utilities representing the restriction of $\preceq$ to $\Delta(X)$ according to (X.7). Let $V^{ut} : \Delta(X) \to \Re$ be defined by $V^{ut}(c) = \min_{u\in\mathcal{U}^{ut}} u \cdot c$; by definition, it is the restriction of V to $\Delta(X)$. Let $K = V^{ut}(\Delta(X))$, and $B(K)$ be the set of functions from $S \to K$. With slight abuse of notation, for $f \in \mathcal{A}$, we use $V^{ut}(f)$ to denote the element of $B(K)$ with $V^{ut}(f)(s) = V^{ut}(f(s))$ for all $s \in S$. By X-A9, the function $I : B(K) \to \Re$ defined by $I(a) = V(f)$ for any f such that $V^{ut}(f) = a$ is well-defined. By definition, $V(f) = I(V^{ut}(f))$.

Lemma X.A.5. *I is concave, continuous, monotonic, and normalised (ie. for all $x \in K$, $I(x^*) = x$, where x^* is the constant function in $B(K)$ taking value x).*

Proof. Normalisation follows from the definition of I. Monotonicity of I is immediate from X-A9. Continuity follows from the continuity of V and the fact that V^{ut} is a quotient map. To establish concavity, we first introduce the following notion. For any $f, g \in \mathcal{A}$, a *state-wise calibrated sequence* from f to g is a sequence $(f_i)_{0\leqslant i\leqslant n}$ with $f_0 = f$, $f_n = g$ such that for every $1 \leqslant i \leqslant n-1$ and every $s \in S$, $\ell(f_{i-1}(s), f_i(s), f_{i+1}(s)) = \frac{1}{2}$, where:

(X.15)
$$\ell(g, f, h) = \begin{cases} \sup\{\alpha \mid \exists g', h' \in \mathcal{A}\ s.t\ g' \sim g, h' \sim h, h'_\alpha g' \sim f\} & \text{if } g \prec f \prec h \\ \sup\{1-\alpha \mid \exists g', h' \in \mathcal{A}\ s.t\ g' \sim g, h' \sim h, h'_\alpha g' \sim f\} & \text{if } g \succ f \succ h \\ \frac{1}{2} & \text{if } g \sim f \sim h \\ 0 & \text{otherwise} \end{cases}$$

We begin with the following claim.

Claim X.5. *For every $f, g \in \mathcal{A}$, for every state-wise calibrated sequence $(f_i)_{0\leqslant i\leqslant n}$ from f to g and every $0 \leqslant i \leqslant n$, $V(f_i) \geqslant \frac{n-i}{n}V(f) + \frac{i}{n}V(g)$.*

Proof. Let $(f_i)_{0\leqslant i\leqslant n}$ be a state-wise calibrated sequence, and consider three successive elements f_{i-1}, f_i, f_{i+1}. We first show that $V(f_i) \geqslant \frac{1}{2}V(f_{i-1}) + \frac{1}{2}V(f_{i+1})$. Since the sequence is state-wise calibrated, for each $s \in S$, $f_i(s) \sim c^s_{\frac{1}{2}} d^s$ for some $c^s, d^s \in \Delta(X)$ with $c^s \sim f_{i-1}(s)$ and $d^s \sim f_{i+1}(s)$. (Notice that the suprema in the definition of ℓ are obtained because $\mathcal{A}$ is compact.) Define $f'_{i-1}, f'_{i+1} \in \mathcal{A}$ by $f'_{i-1}(s) = c^s$ and $f'_{i+1}(s) = d^s$ for all $s \in S$. By X-A9, $f_i \sim f'_{i-1 \frac{1}{2}} f'_{i+1}$. Since V represents $\preceq$ and is concave, $V(f_i) \geqslant \frac{1}{2}V(f'_{i-1}) + \frac{1}{2}V(f'_{i+1}) = \frac{1}{2}V(f_{i-1}) + \frac{1}{2}V(f_{i+1})$, as required.

Applying this fact inductively yields that $V(f_i) \geqslant \frac{1}{i+1}V(f_0) + \frac{i}{i+1}V(f_{i+1})$, and $V(f_i) \geqslant \frac{n-i}{n-i+1}V(f_{i-1}) + \frac{1}{n-i+1}V(f_n)$. It follows (plugging the second inequality into the first) that $V(f_i) \geqslant \frac{n-i}{n}V(f_0) + \frac{i}{n}V(f_n)$, as required. □

To establish the concavity of I, it suffices to show that, for every $a, b \in B(K)$, $I(\frac{1}{2}a + \frac{1}{2}b) \geqslant \frac{1}{2}I(a) + \frac{1}{2}I(b)$. Consider any such a, b with $a \neq b$, and let $f, g \in \mathcal{A}$ be such that $a = V^{ut}(f)$ and $b = V^{ut}(g)$.

Take a sequence of state-wise calibrated sequences $f = f_{0,n}, \ldots, f_{i,n}, \ldots, f_{2n,n} = g$ from f to g; such a sequence can be constructed as in the proof of Theorem X.1. Since $\mathcal{A}$ is compact, there is a subsequence of $f_{n,n}$ that converges as $n \to \infty$. Consider this subsequence of the sequences $\{f_{i,n}\}$ and let $f_{n,n} \to h$. Since $a \neq b$, there exists $s \in S$ such that $f(s) \nsim g(s)$. It follows from the definition of state-wise calibrated sequences and of the function α introduced in the proof of Theorem X.1 (equation (X.13)), that, for every $s \in S$ such that $f(s) \nsim g(s)$, and every n and $0 < i < 2n$:

$$\begin{aligned}\alpha(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s)) =& \frac{\ell(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s))}{1 - \ell(f_{i-1,n}(s), f_{i,n}(s), f_{i+1}(s))} \\ =& 1\end{aligned}$$

Consider any $s \in S$ such that $f(s) \nsim g(s)$ and suppose without loss of generality that $f(s) \prec g(s)$. It follows from the construction used by Kannai (1977, Theorem 2.4) that the restriction of V to $\{f' \in \mathcal{A} \mid f(s) \preceq f' \preceq g(s)\}$ is the limit of a sequence of concave functions V_n, each of which takes value $\frac{\sum_{k=1}^{n} \prod_{i=k}^{2n-1} \alpha(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s))}{\sum_{k=1}^{2n} \prod_{i=k}^{2n-1} \alpha(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s))} V(f(s)) + \frac{\sum_{k=n+1}^{2n} \prod_{i=k}^{2n-1} \alpha(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s))}{\sum_{k=1}^{2n} \prod_{i=k}^{2n-1} \alpha(f_{i-1,n}(s), f_{i,n}(s), f_{i+1,n}(s))} V(g(s))$ on $f_{n,n}(s)$. Since this tends to $\frac{1}{2}V(f(s)) + \frac{1}{2}V(g(s))$ as $n \to \infty$, $V(h(s)) = \frac{1}{2}V(f(s)) + \frac{1}{2}V(g(s))$. Since this holds for every s such that $f(s) \nsim g(s)$, $V^{ut}(h) = \frac{1}{2}a + \frac{1}{2}b$. By Claim X.5, for each sequence n, $V(f_{n,n}) \geqslant \frac{1}{2}V(f) + \frac{1}{2}V(g)$, so by the continuity of V, $V(h) \geqslant \frac{1}{2}V(f) + \frac{1}{2}V(g)$. Hence $I(\frac{1}{2}a + \frac{1}{2}b) \geqslant \frac{1}{2}I(a) + \frac{1}{2}I(b)$, as required. So I is concave.

□

Since $B(K)$ is closed and bounded and I is concave and continuous, for each $a \in B(K)$, there exists an affine functional $\phi : \Re^S \to \Re$ supporting I at a: i.e. such that $\phi(b) \geqslant I(b)$ for all $b \in B(K)$ and $\phi(a) = I(a)$. Each such ϕ can be written as $\phi(b) = x \cdot b + \mu$ for some $x \in \Re^S$, $\mu \in \Re$. We first show that, if such a ϕ supports I at an interior point, then $x(s) \geqslant 0$ for every $s \in S$. Take any $a \in ri(B(K))$ such that ϕ supports I at a, and consider $a_s \in B(K)$ defined by $a_s(s) = a(s) + \epsilon$, $a_s(s') = a(s)$ for $s' \neq s$, where $\epsilon > 0$ such that $a(s) + \epsilon \in K$. By the monotonicity of I, $I(a_s) \geqslant I(a)$; since ϕ supports I at a, we have that $\phi(a_s) = x \cdot a_s + \mu \geqslant I(a_s) \geqslant I(a) = x \cdot a + \mu$, so $x(s).\epsilon \geqslant 0$. Since this holds for every $s \in S$, we have that $x(s) \geqslant 0$ for every $s \in S$. Moreover, since I is non-constant, $x \neq 0$. So each ϕ supporting I at an interior point can be written in the form $\phi(a) = (\bar{a}.p) \cdot a + \bar{b}$ for some $p \in \Delta \subseteq \Re^S$, $\bar{a} \in \Re_{>0}$ and $\bar{b} \in \Re$. Let $\mathcal{P} = cl\left(conv\{(p, \bar{a}, \bar{b}) \in \Delta \times \Re_{>0} \times \Re \mid \exists a \in ri(B(K))\ s.t.(\bar{a}.p) + \bar{b} \text{ supports } I \text{ at } a\}\right)$. By the fact that I is differentiable on a dense subset of $ri(B(K))$, the continuity of the superdifferential mapping, and the continuity of I, $I(a) = \min_{(p,\bar{a},\bar{b}) \in \mathcal{P}} (\bar{a}.p) \cdot a + \bar{b}$ for all $a \in B(K)$. By the fact that supergradients at differentiable points determine the supergradients elsewhere (Rockafellar, 1970, 25.6), $\mathcal{P}$ is tight: no proper closed convex subset of it represents I in this way. For all $x \in K$, since $I(x^*) = x$, $\bar{a}x + b \geqslant x$ for all $(p, \bar{a}, \bar{b}) \in \mathcal{P}$. Moreover, we argue that there exists $p \in \Delta$ such that $(p, 1, 0) \in \mathcal{P}$. To show this, note firstly that for every $x > \min K$, there exists no $(p, a, b) \in \mathcal{P}$ such that $a > 1$ and $a.x + b = x$: if there were such a, b, then for any $y \in K$ with $y < x$, $a.y + b = x - a.(x - y) < y$, contradicting the fact that $\mathcal{P}$ represents I and $I(y^*) = y$. Similarly, for every $x < \max K$, there exists no $(p, a, b) \in \mathcal{P}$ such that $a < 1$ and $a.x + b = x$. Since I is normalised, it follows from the representation of I by $\mathcal{P}$ that there exists $p \in \Delta$ with $(p, 1, 0) \in \mathcal{P}$.

Define $\alpha : \Delta \to \Re_{>0} \times \Re$ by $\alpha(p) = \{(a, b) \in \Re_{>0} \times \Re \mid (p, a, b) \in \mathcal{P}\}$. By the definition of $\mathcal{P}$, and in particular the fact that it is non-empty, closed, convex and tight, α is non-trivial, upper hemicontinuous, convex and tight. By the last two properties of $\mathcal{P}$ mentioned above, α is calibrated (with respect to $\mathcal{U}^{ut}$) and grounded. So α is a non-trivial, tight, grounded, calibrated, upper hemicontinuous, convex function representing $\preceq$ according to (X.11), as required.

It remains to show that $\mathcal{U}^{ut}$ is minimal. This is established by the following lemma.

Lemma X.A.6. *$\mathcal{U}_{\Delta(X)}$ is a minimal tight closed convex representation of the restriction of $\preceq$ to $\Delta(X)$.*

Proof. By Lemma X.A.1, it suffices to show that $V|_{\Delta(X)}$ is a minimal concave representation of the restriction of $\preceq$ to $\Delta(X)$. Suppose that this is not the case, so there exists a concave functional on $\Delta(X)$, v', representing the restriction of $\preceq$ to $\Delta(X)$ and such that $v'(\Delta(X)) = V|_{\Delta(X)}(\Delta(X))$ and $V|_{\Delta(X)}(c) \geqslant v'(c)$ for all $c \in \Delta(X)$ with strict inequality for some c (Kannai, 1977; Debreu, 1976). By Lemma X.A.5, there exists a normalised, monotone, concave functional I such that $V = I \circ V|_{\Delta(X)}$. Hence $V' = I \circ v'$ represents $\preceq$, is concave and is such that $V(f) \geqslant V'(f)$ for all $f \in \mathcal{A}$ with strict inequality for some f, contradicting the minimality of V. So $\mathcal{U}_{\Delta(X)}$ is minimal, as required.

□

Consider finally the uniqueness clause. The uniqueness of $\mathcal{U}^{ut}$ follows from the arguments used to establish the uniqueness clause in Theorem X.1, applied to the restriction to constant acts. Suppose that $\mathcal{U}^{ut}$ and α, and $\mathcal{U}^{ut}$ and α' represent $\preceq$. Then α and α' both represent the same functional $I : B(K) \to \Re$: $I(a) \leqslant I(b)$ iff $\min_{p\in\Delta,\ (\bar{a},\bar{b})\in\alpha(p)} \bar{a}(a \cdot p) + \bar{b} \leqslant \min_{p\in\Delta,\ (\bar{a},\bar{b})\in\alpha(p)} \bar{a}(b \cdot p) + \bar{b}$, for all $a, b \in B(K)$, and similarly for α'. Since α and α' are grounded and calibrated, $I(a) = \min_{p\in\Delta,\ (\bar{a},\bar{b})\in\alpha(p)} \bar{a}(a\cdot p)+\bar{b} = \min_{p\in\Delta,\ (\bar{a},\bar{b})\in\alpha'(p)} \bar{a}(a\cdot p)+\bar{b}$ for all $a \in B(K)$. By the arguments in the proof of Theorem X.1, $\alpha = \alpha'$, as required.

□

Bibliography

Anscombe, F. J. and Aumann, R. J. (1963). A Definition of Subjective Probability. *The Annals of Mathematical Statistics*, 34:199–205.

Arrow, K. J. (1974). Optimal Insurance and Generalized Deductibles. *Scandinavian Actuarial Journal*, 1:1–42.

Bradley, S. (2014). Imprecise Probabilities. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2014 edition.

Cerreia-Vioglio, S. (2009). Maxmin expected utility on a subjective state space: convex preferences under risk. Technical report, Columbia University.

Cerreia-Vioglio, S., Ghirardato, P., Maccheroni, F., Marinacci, M., and Siniscalchi, M. (2011a). Rational preferences under ambiguity. *Economic Theory*, 48(2-3):341–375.

Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., and Montrucchio, L. (2011b). Uncertainty averse preferences. *Journal of Economic Theory*, 146(4):1275–1330.

Chateauneuf, A. and Faro, J. H. (2009). Ambiguity through confidence functions. *J. Math. Econ.*, 45:535–558.

Cook, P. J. and Graham, D. A. (1977). The Demand for Insurance and Protection: The Case of Irreplaceable Commodities. *The Quarterly Journal of Economics*, 91:143–156.

Debreu, G. (1976). Least concave utility functions. *Journal of Mathematical Economics*, 3(2):121–129.

Drèze, J. H. (1987). *Essays on Economic Decisions under Uncertainty*. Cambridge University Press, Cambridge.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quart. J. Econ.*, 75(4):643–669.

Fishburn, P. C. (1970). *Utility Theory for Decision Making*. Wiley, New York.

Galaabaatar, T. and Karni, E. (2013). Subjective expected utility with incomplete preferences. *Econometrica*, 81(1):255–284.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *J. Math. Econ.*, 18(2):141–153.

Hill, B. (2010). An additively separable representation in the Savage framework. *J. Econ. Theory*, 145(5):2044–2054.

Kannai, Y. (1977). Concavifiability and constructions of concave utility functions. *Journal of mathematical Economics*, 4(1):1–56.

Karni, E. (1983). Risk Aversion for State-Dependent Utility Functions: Measurement and Applications. *International Economic Review*, 24:637–647.

Karni, E. (1993a). A Definition of Subjective Probabilities with State-Dependent Preferences. *Econometrica*, 61:187–198.

Karni, E. (1993b). Subjective expected utility theory with state dependent preferences. *J. Econ. Theory*, 60:428–438.

Karni, E. (2011). A theory of Bayesian decision making with action-dependent subjective probabilities. *Economic Theory*, 48(1):125–146.

Karni, E. and Mongin, P. (2000). On the Determination of Subjective Probability by Choices. *Management Science*, 46:233–248.

Karni, E. and Schmeidler, D. (1993). On the uniqueness of subjective probabilities. *Economic Theory*, 3:267–277.

Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51:1021–1032.

Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6):1849–1892.

Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, volume 1. Academic Press, San Diego.

Maccheroni, F. (2002). Maxmin under risk. *Economic Theory*, 19(4):823–831.

Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 74(6):1447–1498.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press, USA.

Nau, R. (2006). The shape of incomplete preferences. *The Annals of Statistics*, 34(5):2430–2448.

Ok, E. A., Ortoleva, P., and Riella, G. (2012). Incomplete Preferences Under Uncertainty: Indecisiveness in Beliefs versus Tastes. *Econometrica*, 80(4):1791–1808.

Rockafellar, R. T. (1970). *Convex Analysis*, volume 28 of *Princeton Mathematics Series*. Princeton University Press.

Savage, L. J. (1954). *The Foundations of Statistics*. Dover, New York.

Schmeidler, D. (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57(3):571–587.

Seidenfeld, T., Schervish, M., and Kadane, J. (1995). A representation of partially ordered preferences. *The Annals of Statistics*, 23(6):2168–2217.

Wakker, P. and Zank, H. (1999). State Dependent Expected Utility for Savage's State Space. *Mathematics of Operations Research*, 24:8–34.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. Chapman and Hall, London.