Journal of Mathematical Economics 65 (2016) 83–103

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Incomplete preferences and confidence ☆

Brian Hill *

GREGHEC, HEC Paris, CNRS, Université Paris-Saclay, 1 rue de la Libération, 78351 Jouy-en-Josas, France

ARTICLE INFO

Article history:
Received 3 March 2016
Received in revised form
24 May 2016
Accepted 26 May 2016
Available online 9 June 2016

Keywords:
Incomplete preferences
Confidence
Multiple priors
Choice under incomplete preferences
Absence of trade

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 representation of confidence in beliefs introduced in Hill (2013), and axiomatise a class of models, differing from each other in the appropriate notion of stakes. The theory naturally suggests two distinct strategies for completing preferences, and hence for choosing in the presence of incompleteness: one that relies only on beliefs in which the decision maker is sufficiently confident, and one that mobilises all beliefs, no matter how little confidence she may have in them. Axiomatic characterisations are given for completion procedures following each of the strategies. Finally, in a market setting, the incorporation of confidence is shown to add an extra friction, beyond the standard implications of non-expected utility models for Pareto optima.

1. Introduction

Incomplete preferences have been increasingly recognised 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, 1986/2002). 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, 1986/2002; Masatlioglu and Ok, 2005). Incomplete preferences naturally arise in 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.

The traditional approach to modelling incomplete preferences proceeds, roughly speaking, by dropping the completeness axiom whilst retaining the other standard axioms, and replacing the single function or measure in the relevant model by a set. For instance, in decision under uncertainty, the benchmark unanimity multi-prior model proposed by Bewley (1986/2002) retains all standard Anscombe and 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 colour of the next ball drawn from an urn containing only black and white balls, as shown in Fig. 1. For simplicity, suppose that the bets are given in dollars and the decision maker has linear utility.[1] 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 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

☆ I would like to thank Alain Chateauneuf, Michéle Cohen, Itzhak Gilboa, Edi Karni, Philippe Mongin, Jean-Marc Tallon, editors and reviewers, and audiences at the Paris workshop on Incomplete Preference (2011) and D-TEA 2011 for valuable discussion, and very helpful comments. The author gratefully acknowledges support from the ANR program DUSUCA (ANR-14-CE29-0003-01).

* Fax: +33 1 39 67 71 09.
E-mail address: hill@hec.fr.

[1] 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; e.g. f is the mixture $\frac{1}{100\,000}g + \frac{99\,999}{100\,000}0$.

http://dx.doi.org/10.1016/j.jmateco.2016.05.007
0304-4068/

	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 $\times n$

Fig. 1. Bets ('M' stands for 'million').

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 (e.g. 10%)[2]—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 agreement among a few members of the 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 traditional 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.

The aim of this paper is to propose a model of decision under uncertainty that, whilst deviating as little as possible from standard models of incomplete preference, incorporates the decision maker's confidence in her beliefs. Inspired by the above considerations, it seems that an appropriate model should adhere to the following maxim: one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form a preference does not match up to the stakes involved in the choice. We develop such a model, drawing on existing research on confidence in belief and its role in decision making, and in particular on the concepts introduced in Hill (2013). Like the standard Bewley model, we focus on indeterminacy of preferences that is driven solely by the decision maker's beliefs, tacitly assuming that she is fully confident in her utilities.

As concerns behavioural properties, note that in the context of incomplete preferences, independence applied to the preference $f \succ 0$ and the acts g and 0 (Fig. 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 ($g \succ 0$). 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.[3] 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.

Another central contribution of the paper is to identify some interesting consequences of the incorporation of confidence 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 characterisation 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 mobilise 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', 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.

Finally, a standard interpretation of indeterminacy of preferences in market settings (dating back at least to Bewley, 1986/2002) is in terms of reluctance to trade, and it is natural to ask what implications the incorporation of confidence into models of incomplete preference has 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 2. The model is formally stated in Section 3.1, and the representation result is given in Sections 3.2–3.3. Section 3.4 contains a comparative statics analysis. In Section 4, we consider the question of how to complete one's incomplete preferences. In Section 5, we consider the consequences of the model in markets under uncertainty. Related literature is discussed in Section 6. Proofs of all results and other material are to be found in the Appendices.

2 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.

3 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 3.2 for further discussion.

2. General preliminaries

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 $\Delta(X)$. For technical convenience, we shall focus on unbounded utilities, where a utility function $u : \Delta(X) \to \Re$ is unbounded if $u(\Delta(X)) = \Re$.

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. For a preference relation $\preceq$, let $\preceq_{\Delta(X)}$ be the associated *risk preferences*; that is, the restriction of $\preceq$ to the set of constant acts $\Delta(X)$. $\mathcal{P}_{\Delta(X)}$ is the set of preference relations on $\Delta(X)$.

For any $\preceq_{\Delta(X)} \in \mathcal{P}_{\Delta(X)}$, a pair of acts $(f, g) \in \mathcal{A} \times \mathcal{A}$ constitutes a *non-trivial choice* if $f(s) \nsim_{\Delta(X)} g(s)$ for some $s \in S$. $(\mathcal{A} \times \mathcal{A})^{nt}_{\preceq_{\Delta(X)}} \subset \mathcal{A} \times \mathcal{A}$ is the set of non-trivial choices. Moreover, for any pair $(f, g) \in \mathcal{A} \times \mathcal{A}$, let $\widehat{(f, g)}_{\preceq_{\Delta(X)}} = \{(f', g') \in \mathcal{A} \times \mathcal{A} \mid \exists \alpha \in (0, 1], h \in \mathcal{A}\ s.t.\ f'(s) \sim_{\Delta(X)} (f_\alpha h)(s)$ and $g'(s) \sim_{\Delta(X)} (g_\alpha h)(s)\ \forall s \in S$, or $f(s) \sim_{\Delta(X)} (f'_\alpha h)(s)$ and $g(s) \sim_{\Delta(X)} (g'_\alpha h)(s)\ \forall s \in S\}$. There is a sense in which the choice between f and g and the choice between $f_\alpha h$ and $g_\alpha h$ are the 'same' choice; we will say that these two choices are *versions* of each other. $\widehat{(f, g)}_{\preceq_{\Delta(X)}}$ is thus the set of all pairs which are versions of the choice between f and g. The standard independence axiom implies that preferences are uniform on such sets (for all $(f', g') \in \widehat{(f, g)}_{\preceq_{\Delta(X)}}$, $f' \preceq g'$ iff $f \preceq g$); as we shall see, this will not be assumed here. To ease notation, we omit the subscript $\preceq_{\Delta(X)}$ from both notions when it is clear from the context.

Since $(\mathcal{A} \times \mathcal{A})^{nt}$ and $\widehat{(f, g)}$ are entirely determined by $\preceq_{\Delta(X)}$, once one knows the utility function representing $\preceq_{\Delta(X)}$, these sets can be fully defined. In practice, there are several model-free methods for measuring utilities (for example, Wakker and Deneffe, 1996; Abdellaoui, 2000), which are independent of the treatment of uncertainty and hence can be applied in the context of this model. Any such method could be used by an analyst to obtain these sets.

2.2. Stakes

Under the maxim proposed in the Introduction, the stakes involved in a choice between two options may have implications for the preferences over them. One can imagine several different notions of what counts as the stakes in a choice; Fig. 2 gives some examples. We shall be non-committal about the 'proper' notion of stakes, and use an abstract representation here. Each notion of stakes specifies whether a choice involves higher or lower stakes than another, often on the basis of the utility profiles of the acts (see Fig. 2). Accordingly, whether a given choice involves higher or lower stakes depends on the decision maker's utility function and hence on his risk preferences, $\preceq_{\Delta(X)}$. So we model the *notion of stakes* by a function $\leqq_\bullet$ that, given $\preceq_{\Delta(X)} \in \mathcal{P}_{\Delta(X)}$, yields a *stakes relation* $\leqq_{\preceq_{\Delta(X)}}$, that is, a relation on the set of non-trivial choices (pairs of acts in $(\mathcal{A} \times \mathcal{A})^{nt}$).[4] This relation is interpreted as follows: $(f, g) \leqq_{\preceq_{\Delta(X)}} (f', g')$ 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'.[5] $\equiv_{\preceq_{\Delta(X)}}$ and $<_{\preceq_{\Delta(X)}}$ are the symmetric and asymmetric components of $\leqq_{\preceq_{\Delta(X)}}$, defined in the standard way. We shall be interested in notions of stakes satisfying the following basic properties, for all $\preceq_{\Delta(X)}$.

(Weak Order) $\leqq_{\preceq_{\Delta(X)}}$ is reflexive, transitive and complete.

(Symmetry) For all $(f, g) \in (\mathcal{A} \times \mathcal{A})^{nt}$, $(f, g) \equiv_{\preceq_{\Delta(X)}} (g, f)$.

(Extensionality) For all $(f, g), (f', g') \in (\mathcal{A} \times \mathcal{A})^{nt}$, if $f(s) \sim_{\Delta(X)} f'(s)$ and $g(s) \sim_{\Delta(X)} g'(s)$ for all $s \in S$, then $(f, g) \equiv_{\preceq_{\Delta(X)}} (f', g')$.

(Continuity) For all $(f, g), (f', g') \in (\mathcal{A} \times \mathcal{A})^{nt}$ and $h \in \mathcal{A}$, the sets $\{(\alpha, \beta) \in [0, 1]^2 | (f_\alpha h, g_\beta h) \geqq_{\preceq_{\Delta(X)}} (f', g')\}$ and $\{(\alpha, \beta) \in [0, 1]^2 | (f_\alpha h, g_\beta h) \leqq_{\preceq_{\Delta(X)}} (f', g')\}$ are closed in $\{(\alpha, \beta) \in [0, 1]^2 | (f_\alpha h, g_\beta h) \in (\mathcal{A} \times \mathcal{A})^{nt}\}$.

(Richness) For all $(f, g), (\bar{f}, \bar{g}) \in (\mathcal{A} \times \mathcal{A})^{nt}$, there exists $(f', g'), (f'', g'') \in \widehat{(f, g)}$ such that $(f', g') \leqq_{\preceq_{\Delta(X)}} (\bar{f}, \bar{g}) \leqq_{\preceq_{\Delta(X)}} (f'', g'')$.

Weak order states that the non-trivial 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.[6] Symmetry states that the stakes involved in a choice do not depend on the order in which the alternatives are presented, and deserves no further discussion.[7]

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. Since this aspect of standard practice is not the focus of the current paper, we shall follow it here and assume extensionality, which essentially says the same thing for stakes.

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 there exists a version of every choice that involves stakes as far up or down the stakes order as desired. As noted in Section 2.1, there is a sense

4 So, formally, a notion of stakes is a function $\leqq_\bullet : \mathcal{P}_{\Delta(X)} \to (\mathcal{A} \times \mathcal{A}) \times (\mathcal{A} \times \mathcal{A})$ whose image is contained in $(\mathcal{A} \times \mathcal{A})^{nt} \times (\mathcal{A} \times \mathcal{A})^{nt}$. Note that it naturally associates a stakes relation to each utility function.

5 Since trivial choices—such as the choice between an act and itself—can only be qualified as choices in a technical sense of the term, it seems inappropriate to talk of them as having stakes; for this reason, they are absent from the domain of the stakes relation. Nevertheless, the main results (with relevant definitions modified accordingly) continue to apply if stakes relations assign a position in the stakes order to trivial choices.

6 See Hill (2012, Section 3.1) for a discussion of the possibility of weakening the completeness assumption in a different but related framework.

7 As shall be evident in the discussion below, symmetry does not rule out dependence on properties of the acts other than the order of presentation, such as whether one of the acts is the status quo.

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 maximum of the negation of the expected utility of the part of the act taking values below some threshold, calculated using a given probability measure, taken over f or g

(iv) the maximum absolute value of the difference between the utility of $f(s)$ and the utility of $g(s)$, taken over $s \in S$

(v) the maximum absolute value of the difference between the utility of $f(s)$ and the utility of $g(s)$, taken at the $s \in S$ where this difference is non-zero for which the minimum utility out of $f(s)$ and $g(s)$ is lowest.

Fig. 2. Some notions of stakes.

in which the choices between f and g, and between $f_\alpha h$ and $g_\alpha h$ are versions of the 'same' choice. Nevertheless, the stakes involved in these two choices may differ; to that extent, the latter choice, for instance, can be thought of as a version of the former one at the stakes level corresponding to $(f_\alpha h, g_\alpha h)$. Using such versions, one can thus 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 prefers f' to g' for some version $(f', g') \in \widehat{(f, g)}$ involving that level of stakes. Richness simply states that for any non-trivial choice and stakes level, there is a version of the choice that has stakes above the level in question, and there is a version that has stakes below that level. The intuition is that the mixing with a third act involved in the definition of the versions of a choice, $\widehat{(f, g)}$, can affect 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.

To get an idea of the mildness of these properties, note that all the notions of stakes in Fig. 2 satisfy them. To ease notation, we omit the subscript $\preceq_{\Delta(X)}$ from the stakes relation $\leqq_{\preceq_{\Delta(X)}}$ when it is clear from the context.

In this paper, we adopt the perspective presented in Hill (2013), according to which the notion of stakes is an objective feature of the decision model: representations involving different notions of stakes are considered to constitute different decision models (in much the same way as representations with different functional forms are usually treated as different models).[8] We thus assume throughout Sections 3 and 4 a notion of stakes $\leqq_\bullet$ satisfying the five properties above. The analysis and results hold for any notion satisfying these properties: the specification of a particular notion of stakes (such as any of the notions in Fig. 2) will yield axioms and results for representations involving that notion. As such, the results below 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 involve stakes, different notions of stakes will correspond to different behavioural properties, and hence it is possible in principle to tell whether the decision maker is using a given notion or not. See Hill (2016) for a representation in which the stakes (over general, rather than only two-element menus) are endogenously derived in a related decision model.

Remark 1. Rather than assuming a notion of stakes yielding a stakes relation defined on pairs of acts, it could have been defined on triples in $\mathcal{A} \times \mathcal{A} \times \Gamma$, where Γ can have several interpretations. Γ could be understood as a set of context indices; hence dependence of the stakes on the context can be accommodated. Alternatively, Γ could be interpreted as the status quo, if there is one; one can thus capture dependence of the stakes on the status quo. It is straightforward to adapt the properties above to such notions of stakes; corresponding modifications to the axioms below yield similar results where the stakes may depend on factors other than the two acts on offer.

2.3. Confidence ranking and cautiousness coefficient

We recall two notions that were introduced in Hill (2013).

Definition 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 *balanced* if, for every $\mathcal{C}_1, \mathcal{C}_2 \in \Xi$ with $\mathcal{C}_1 \subset \mathcal{C}_2$, $\mathcal{C}_1 \cap ri(\mathcal{C}_2) \neq \varnothing$.[9]

Confidence rankings represent decision makers' confidence in their beliefs. A set in the confidence ranking is interpreted as corresponding to a level of confidence. A probability judgement[10] that applies under every probability measure in the set is one that the decision maker holds to the corresponding level of confidence. Larger sets correspond to higher levels of confidence; this translates the fact that one holds fewer probability judgements (or beliefs) with 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

[8] The reader is referred to Section 6 and especially Hill (2013) for a discussion of how different decision models in the same family can be obtained by varying decision rules and notions of stakes.

[9] For a set X, $\overline{X}$ is the closure of X and $ri(X)$ is its relative interior. Note that the union of a nested family of convex sets is convex.

[10] By probability judgement, we mean a statement concerning probabilities, such as 'the probability of event A is greater than p'.

as corresponding to a level in the group's hierarchical structure (e.g. in a country, one level will contain cabinet ministers, another will contain members of the government, another all elected representatives, and so on). A probability judgement 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 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 is in $[0.3, 0.7]$ and then confident only that the probability is in $[0.1, 0.9]$ at the 'next' confidence level up. Balancedness—the sole property not already present in Hill (2013)—guarantees that, whilst precise probability assignments held at lower confidence levels may be revoked at higher confidence levels, this will not happen in a lopsided way. For example, if one is confident that the probability of an event is 0.5 at some confidence level, balancedness allows one to be confident only that the probability is in $[0.45, 0.55]$ at a higher level, but it does not permit one to be confident only that it is in an interval such as $[0.5, 0.6]$ (where 0.5 is on the boundary). This property of the confidence ranking is a direct consequence of retaining the essence of the standard independence property (Section 3.2), when weak preferences are taken as primitive.

The second notion required is that of a *cautiousness coefficient* for a confidence ranking Ξ, under a utility function u, which is defined to be a function $D : \mathcal{A} \times \mathcal{A} \to \Xi$ that is surjective on $(\mathcal{A} \times \mathcal{A})^{nt}$—for each $\mathcal{C} \in \Xi$, there exists $(f, g) \in (\mathcal{A} \times \mathcal{A})^{nt}$ with $D((f, g)) = \mathcal{C}$—and preserves $\leqq$—for all $(f, g), (f', g') \in (\mathcal{A} \times \mathcal{A})^{nt}$, if $(f, g) \leqq (f', g')$, then $D((f, g)) \subseteq D((f', g'))$. The cautiousness coefficient represents the decision maker's attitude to choosing in the absence of confidence. It assigns to any pair of acts the level of confidence that is required in beliefs for them to play a role in the choice between the acts. This level of confidence corresponds to the appropriate set of probability measures in the confidence ranking. Preservation of the stakes relation 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.[11] Since in this paper stakes are assigned to choices (pairs of acts), it is natural to take this to be the domain of the cautiousness coefficient. (Hill, 2013 uses a notion of stakes, and accordingly a cautiousness coefficient, defined on acts; see Sections 4.3 and 6.) 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 a choice, for which it is the relevant set.

As suggested, the confidence ranking and the cautiousness coefficient are (like the utility function) subjective elements in the model, representing the decision maker's attitudes—specifically, beliefs and (certain) tastes respectively. This separation of attitudes, which is defended formally in Section 3.4, allows a clean connection between these technical notions and the everyday concept of confidence in beliefs. In particular, the informal English-language concept of confidence in beliefs denotes an aspect of an individual's beliefs, separate from her tastes—and the same goes for the formalisation proposed here, namely the confidence ranking. The reader is referred to Hill (2013) for further discussion and defence.

[11] D can assign any confidence level to a trivial choice (a pair of acts not in $(\mathcal{A} \times \mathcal{A})^{nt}$): since the choice is trivial, the assignment will have no effect on preferences. Since these are uninteresting, we shall henceforth focus the discussion on non-trivial choices.

3. A theory of incomplete preferences and confidence

In this section, we axiomatise and analyse a representation of incomplete preferences. Throughout this section and the next one, we assume a fixed notion of stakes $\leqq_\bullet$ satisfying the properties in Section 2.2.

3.1. Representation

The representation of preferences that we shall consider is of the following form: for all acts f and g, $f \preceq g$ if and only if

$$\sum_{s\in S} u(f(s)).p(s) \leqslant \sum_{s\in S} u(g(s)).p(s) \quad \text{for all } p \in D((f, g)) \tag{1}$$

where u is a utility function on $\Delta(X)$ and D is a cautiousness coefficient for a confidence ranking Ξ, under u. This representation embodies the guiding maxim stated in the Introduction: one's preferences are indeterminate when and only when one's confidence in the beliefs needed to form a preference does not match up to the stakes involved in the choice. To see why, 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. As explained in Section 2.3, it specifies the confidence level associated to those stakes. So a decision maker is confident enough in a probability judgement for it to play a role in the decision only if it holds for all probability measures in $D((f, g))$. Under representation (1), g is weakly preferred to f if and only if, based only on such 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 does not hold the beliefs needed to conclude in favour of one of them to the level of confidence required by the stakes involved in the decision. Representation (1) is thus a faithful formal rendition of the aforementioned maxim.

3.2. Axioms

Consider the following axioms on $\preceq$.

Axiom A1 (*C-Completeness*). For all $c, d \in \Delta(X)$, $c \asymp d$.

Axiom A2 (*Reflexivity and Non-Degeneracy*). $\preceq$ is reflexive and non-degenerate.

Axiom A3 (*Stakes Transitivity*). For all $f, g, h \in \mathcal{A}$, $(f', g') \in \widehat{(f, g)}$ and $(g'', h'') \in \widehat{(g, h)}$ such that $(f, h) \leqq (f', g')$ or $f(s) \sim g(s)$ for all $s \in S$, and $(f, h) \leqq (g'', h'')$ or $g(s) \sim h(s)$ for all $s \in S$, if $f' \preceq g'$ and $g'' \preceq h''$, then $f \preceq h$.

Axiom 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 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 A6 (*Consistency*). For all $f, g \in \mathcal{A}$ and $(f', g') \in \widehat{(f, g)}$ such that $(f', g') \leqq (f, g)$, if $f \asymp g$, then $f' \asymp g'$.

Axiom 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 A8 (*Continuity in Stakes*). For all $(f, g) \in (\mathcal{A} \times \mathcal{A})^{nt}$, $h \in \mathcal{A}$ 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 $(f', (g_\beta h)') \in \widehat{(f, g_\beta h)}$ such that $(f', (g_\beta h)') > (f, g)$ and $f' \succeq (g_\beta h)'$.

Axiom A9 (*Unboundedness*). For every $c, d \in \Delta(X)$ with $c \prec d$, there exists $e, e' \in \Delta(X)$ such that $d_{\frac{1}{2}}e \preceq c \prec d \preceq c_{\frac{1}{2}}e'$.

Reflexivity and Non-degeneracy (A2) and Monotonicity (A5) are standard and require no further comment. Continuity (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 (A1), which corresponds to an axiom first introduced by Bewley (1986/2002) and proposed under this name 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 (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 repeatedly) 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 the choices are trivial, 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. (Recall from Sections 2.1 and 2.2 that preferences over acts f and g at a given stakes level are fleshed out formally in terms of preferences over versions of that choice—elements of $\widehat{(f, g)}$—having those stakes.) 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 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.[12] In this sense, this is a particularly mild weakening of transitivity. Note finally that, in the presence of A4, stakes transitivity is equivalent to transitivity whenever preferences are complete.

A similar situation holds for Pure Independence (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 the Introduction. 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 (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 in a choice between f and g are determinate, then preferences will be determinate in any version of the choice between f and 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 that translates the idea that the higher the stakes, the more confidence is needed to take the choice. This is a fully behavioural axiom, which is in principle testable by, for example, comparing preferences at different stakes levels. (Of course, the other axioms are as behavioural as their standard counterparts.)

Continuity in Stakes (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 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 that $\leqq$ has no minimal elements (as is the case for several of the notions of stakes given in Fig. 2). Unboundedness (A9) is a standard axiom guaranteeing the unboundedness of utility. As mentioned previously, we assume unbounded utilities for technical convenience: modulo slight but clumsy reformulations of some of the axioms and definitions, the results go through in the absence of this assumption.

3.3. Result

The preceding axioms characterise the following representation of preferences.

Theorem 1. *Let $\preceq$ be a binary relation on $\mathcal{A}$. The following are equivalent.*

(i) *$\preceq$ satisfies A1–A9.*
(ii) *There exists an unbounded affine utility function $u : \Delta(X) \to \Re$, a balanced continuous confidence ranking Ξ and a cautiousness coefficient D for Ξ under u such that, for all $f, g \in \mathcal{A}$, $f \preceq g$ if and only if*

$$\sum_{s \in S} u(f(s)).p(s) \leqslant \sum_{s \in S} u(g(s)).p(s) \quad \text{for all } p \in D((f, g)). \quad (1)$$

Furthermore, Ξ is unique, D is unique on $(\mathcal{A} \times \mathcal{A})^{nt}$, and u is unique up to positive affine transformation.

[12] 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.

Note that it follows from this theorem that, if $\preceq$ satisfies A1–A9 and is complete, then $\Xi = \{\{p\}\}$ for some probability $p \in \Delta(S)$, and $D((f,g)) = \{p\}$ for all $(f,g) \in \mathcal{A} \times \mathcal{A}$. In other words, adding completeness to the other axioms yields a standard subjective expected utility representation.

3.4. 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 characterisation of decisiveness in this model, the analysis will also corroborate the interpretations of the confidence ranking and cautiousness coefficient proposed in Section 2.3.

Under models of incomplete preferences, if decision maker 2 weakly prefers f to g whenever decision maker 1 does, this is an indication that 2 is less prone to indeterminacy of preference than 1. This insight inspires the following standard definition of decisiveness: $\preceq^1$ is *less decisive* than $\preceq^2$ if $\preceq^1 \subseteq \preceq^2$.[13]

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 1. *Let $\preceq^1$ and $\preceq^2$ satisfy Axioms A1–A9, 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 $(f,g) \in (\mathcal{A} \times \mathcal{A})^{nt}$.*

Besides the utility function, 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 confidence ranking Ξ_2 with a range of sets of probability measures, and an appropriate cautiousness coefficient, sending different pairs 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 1. To tease them apart, let us introduce the notion of confidence in preferences.

Definition 2. Let $\preceq$ satisfy Axioms A1–A9. The *confidence-in-preferences* relation $\preccurlyeq$ on $(\mathcal{A} \times \mathcal{A})^{nt}$ is defined as follows: for any $f, g, \bar{f}, \bar{g} \in \mathcal{A}$, $(f,g) \preccurlyeq (\bar{f}, \bar{g})$ iff, for all $(f', g') \in \widehat{(f,g)}$ and $(\bar{f}', \bar{g}') \in \widehat{(\bar{f}, \bar{g})}$ such that $(f', g') \equiv (\bar{f}', \bar{g}')$:

$$f' \succeq g' \Rightarrow \bar{f}' \succeq \bar{g}'.$$

Definition 2 relies on the observation that, if a decision maker prefers $\bar{f}$ to $\bar{g}$ 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 $\bar{f}$ over $\bar{g}$ than in her preference for f over g*.[14] 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 1 and 2 are *confidence equivalent* if they have the same confidence-in-preferences relation: $\preccurlyeq^1 = \preccurlyeq^2$.

Proposition 2. *Let $\preceq^1$ and $\preceq^2$ satisfy Axioms A1–A9, 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 $\Xi_1 = \Xi_2$.*

So a decision maker's confidence in her preferences is entirely determined by her utility function and her confidence ranking. Since one would expect a decision maker's confidence in her preferences to be fully determined by her tastes for consequences (utilities) and her confidence in her beliefs, this corroborates the interpretation of the confidence ranking as representing confidence in beliefs.

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. The following corollary of Propositions 1 and 2 makes this explicit.

Corollary 1. *Let $\preceq^1$ and $\preceq^2$ satisfy Axioms A1–A9, 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 $(f,g) \in (\mathcal{A} \times \mathcal{A})^{nt}$.*

In summary, for decision makers with the same confidence in preferences, differences in decisiveness are completely characterised 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.

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, 1986/2002); they may postpone the decision, if possible (Danan, 2003a; Kopylov, 2009); more generally, they may take a deferral option, if one is present (Hill, 2016). However, in many situations, such 'choice-avoidance mechanisms' are unavailable

[13] Containment of the preference relations is equivalent to saying that, for all $f, g \in \mathcal{A}$, if $f \succeq^1 g$ then $f \succeq^2 g$.

[14] Recall from Sections 2.1 and 2.2 that talk of preferences at different stakes levels is spelt out formally in terms of preferences over appropriate versions of the choice. Moreover, for reasons similar to those mentioned in footnote 5, it seems unnatural to speak of—and define—confidence in preferences concerning a trivial choice. Nevertheless, the notion can be simply extended to encompass trivial choices (by stipulating that one has maximal confidence in preferences over them, for example) without affecting the results below.

and the decision maker will have to make a choice. 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 the interpretations of incomplete preferences mentioned in the Introduction, though the form it takes may depend 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 behavioural 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 behaviour 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.[15]

Consider a decision maker who is forced to choose between options over which her preferences are indeterminate. 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 appropriate level of confidence given the stakes involved. 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 (1), the decision maker has to employ a different decision rule, involving considerations of caution or an element of random choice, for example. 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 that, whilst the decision maker would not decide on a hunch—a belief in which she has limited confidence—if she could avoid it (by deferring the decision for instance), when she is forced to decide she may as well mobilise all of her beliefs—even hunches. Given that she is relying on more beliefs, the decision maker may be able to form a determinate preference using the 'unanimity' rule employed in representation (1); if not, she will require a different decision rule.

We suggest that this distinction corresponds to an important difference between possible reactions to forced choice, under all of the interpretations cited in the Introduction. 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 at 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: this is tantamount to adopting a strategy that goes on hunches. Finally, the distinction takes a particularly simple form under the group interpretation of incompleteness: in cases of disagreement within 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.

In this section, we provide an axiomatic analysis of choice on the basis of incomplete preferences, proposing two general procedures that respect confidence, and showing, for each one, the behavioural difference with the corresponding procedure that goes on hunches.

4.1. Preference completions: framework and basic properties

Formally, 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 2 and 3. The issue of choice on the basis of incomplete preferences can be tackled by considering axioms on the completed preference relation and the relationship with $\preceq$, such as the following.

Axiom C1 (*Forced Choice*)**.** $\preceq^c$ is complete.

Axiom B1 (*C-Consistency*)**.** For all $x, y \in \Delta(X)$, $x \preceq y$ if and only if $x \preceq^c y$.

Axiom B2 (*Consistency*)**.** For all $f, g \in \mathcal{A}$, if $g \preceq f$, then $g \preceq^c f$.

Forced Choice (C1) states that the completed preferences are indeed complete; it translates the fact that the decision maker must choose. C-Consistency (B1) just says that the preference orders coincide on constant acts. This is natural, given that incompleteness of $\preceq$ is driven by beliefs, and indeed only occurs for comparisons involving non-constant acts. Consistency (B2), which was first introduced by Gilboa et al. (2010), states that $\preceq^c$ simply completes $\preceq$ without reversing any preferences. These three axioms are the minimal requirement one could ask of a completion of $\preceq^c$; indeed, when we speak of a *completion* of $\preceq$, we shall henceforth mean a preference relation satisfying B1, B2 and C1 with respect to $\preceq$.[16]

4.2. Simple rules for preference completion

To introduce perhaps the most immediate sort of completion rule, consider the following two axioms.

Axiom B3 (*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 $f' \succeq c'$ for some $(f', c') \in \widehat{(f, c)}$ with $(f', c') \equiv (f, g)$, and $g' \not\succeq d'$ for all $(g', d') \in \widehat{(g, d)}$ with $(g', d') \equiv (f, g)$.

Axiom B3^{S-N} (*Stakes-Neutral 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 $f' \succeq c'$ for some $(f', c') \in \widehat{(f, c)}$ and $g' \not\succeq d'$ for all $(g', d') \in \widehat{(g, d)}$.

Incomplete preferences provide a crude indication of the relative worth of acts for the decision maker. One way of getting a more refined judgement is by comparing them with, or 'benchmarking' them against constant acts—that is, by considering which constant acts they are preferred to. 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

[15] See Hill (2012) for further discussion of this interpretation of incompleteness, and Hill (2016) for a general treatment of deferral in the context of decision under uncertainty.

[16] This terminology was introduced in Danan et al. (forthcoming).

a particular constant act (e.g. a sure \$5), whilst not having a determinate preference for g over an inferior constant act (e.g. a sure \$4). The axioms both demand that in precisely these sorts of 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 better than the other act in the comparison with constant acts (according to the initial, incomplete, preference relation). The difference between the axioms is that whilst Benchmark on Certainty (B3) employs stakes-corrected comparisons with constant acts—only versions of the choice between an act and a given constant act at the appropriate stakes level are involved—Stakes-neutral Benchmark on Certainty ($B3^{S-N}$) adopts stakes-neutral comparisons—all versions of the choice are considered, no matter the stakes level. Note that, when $\preceq$ is represented according to (1), both axioms imply C-Consistency (B1) and Consistency (B2).

As the following results show, these two axioms characterise behaviourally the difference between the two strategies described above.

Theorem 2. *Let $\preceq$ satisfy A1–A9, and be represented according to (1) by (u, Ξ, D). Then*

(i) *$\preceq^c$ is a completion of $\preceq$ satisfying B3 if and only if, for all $f, g \in \mathcal{A}$,*

$$f \preceq^c g \quad \text{iff} \quad \min_{p\in D((f,g))} \sum_{s\in S} u(f(s)).p(s) \leqslant \min_{p\in D((f,g))} \sum_{s\in S} u(g(s)).p(s) \tag{2}$$

(ii) *$\preceq^c$ is a completion of $\preceq$ satisfying $B3^{S-N}$ if and only if, for all $f, g \in \mathcal{A}$,*

$$f \preceq^c g \quad \text{iff} \quad \min_{p\in \bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s) \leqslant \min_{p\in \bigcap_{C\in\Xi} C} \sum_{s\in S} u(g(s)).p(s). \tag{3}$$

Both Benchmark on Certainty (B3) and Stakes-neutral Benchmark on Certainty ($B3^{S-N}$) characterise 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, and hence the beliefs used. Benchmark on Certainty (B3) implies that the decision maker uses only the beliefs that she holds with sufficient confidence given the stakes (see the interpretation of $D((f, g))$ in Sections 2.3 and 3.1). A decision maker satisfying this axiom employs a strategy that respects confidence to decide when forced to choose. By contrast, Stakes-neutral Benchmark on Certainty ($B3^{S-N}$) 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. A decision maker satisfying this axiom thus relies on all of her beliefs when forced to choose: she employs a strategy that goes on hunches.

Hence, among cautious procedures for 'completing' incomplete preferences, the difference between Benchmark on Certainty and its stakes-neutral counterpart characterises precisely the difference between a strategy that respects confidence and 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 notion of confidence ranking. To the knowledge of the author, this is the first model in the literature capable of capturing this distinction.

Note finally that representation (2) is by no means the only completion procedure that respects confidence: one could imagine analogous representations with different decision rules such as the α-maxmin EU rule (Ghirardato et al., 2004) or a random decision rule in the style of Gul and Pesendorfer (2006) and Karni and Safra (2014). For each such representation, there will be a 'sister' version, analogous to representation (3), adopting a strategy that goes on hunches. However, the family of completion procedures respecting confidence goes beyond even these examples, as we shall now see.

4.3. Transitivity of the completion and stakes on acts

As the following example illustrates, the completion procedure respecting confidence characterised above (representation (2)) has an important drawback: it may lead to violations of transitivity.

Example 1. Consider a state space consisting of two states, $S = \{s, t\}$ and a real outcome space, $X = \Re$. Suppose that the notion of stakes is $\leqq_{\bullet}$, where $(f, g) \leqq_{\preceq_{\Delta(X)}} (f', g')$ iff $\min_{\preceq_{\Delta(X)}}\{f(s), f(t), g(s), g(t)\} \succeq_{\Delta(X)} \min_{\preceq_{\Delta(X)}}\{f'(s), f'(t), g'(s), g'(t)\}$. $\preceq$ is represented according to (1) with u the identity function and $\Xi = \{\{p : p(s) \in [0.6 - \epsilon, 0.6 + \epsilon]\} : \epsilon \in [0, 0.4]\}$. Since Ξ is naturally parametrised by $\epsilon \in [0, 0.4]$, D can be defined as a function into $[0, 0.4]$. Using this formulation, D is defined by: $D(f, g) = 0$ if $\min\{f(s), f(t), g(s), g(t)\} > 0$, $D(f, g) = -\frac{1}{10000}\min\{f(s), f(t), g(s), g(t)\}$ if $\min\{f(s), f(t), g(s), g(t)\} \in [-4000, 0]$, and $D(f, g) = 0.4$ if $\min\{f(s), f(t), g(s), g(t)\} < -4000$.

Consider the acts f, g, h defined as follows: $f(s) = f(t) = 500$; $g(s) = 4100$, $g(t) = -2000$ and $h(s) = 1000$, $h(t) = 0$. If the completed preferences are represented according to (2), then, by simple calculation, we have $f \succ^c g$ and $g \succ^c h$ (since $D((f, g)) = D((g, h)) = \{p : p(s) \in [0.4, 0.8]\}$), whereas $h \succ^c f$ (since $D((f, g)) = D((f, h)) = \{p : p(s) = 0.6\}$), violating transitivity.

Does this discredit every completion procedure that respects confidence, or is this problem specific to the simple rule encapsulated in representation (2)? Such violations seem to stem from the fact that the notion of stakes used here is defined on (binary) choices—that is pairs of acts—so different sets of probabilities may be used in the evaluation of the same act f in the context of different pairwise comparisons. Indeed, we now present results showing that this problem can be avoided by completing preferences in such a way that stakes can be thought of as being assigned to the acts themselves. Moreover, the behavioural distinction between strategies that respect confidence and those that go on hunches is retained for such completions.

To present the results, we first state a stakes-neutral version of the Consistency axiom, which, analogously with the Stakes-neutral Benchmark on Certainty $B3^{S-N}$, involves preferences over any version of the choice, no matter the stakes.

Axiom $B2^{S-N}$ (*Stakes-Neutral Consistency*). For all $f, g \in \mathcal{A}$, if there exists $(f', g') \in \widehat{(f, g)}$ with $g \preceq f$, then $g \preceq^c f$.

Note that this axiom implies standard Consistency (B2).

We shall also require the following terminology. For any $f \in \mathcal{A}$, let $\underline{c_f} \in \Delta(X)$ be a $\preceq$-minimal element such that, for all $c \in \Delta(X)$, if $f \succeq c$, then $\underline{c_f} \succeq c$. Similarly, let $\overline{c_f} \in \Delta(X)$ be a $\preceq$-maximal element such that, for all $c \in \Delta(X)$, if $f \preceq c$, then $\overline{c_f} \preceq c$. $\underline{c_f}$ can be thought of as a lower certainty equivalent of f—for any completion of $\preceq$, the certainty equivalent will be weakly preferred to $\underline{c_f}$; similarly, $\overline{c_f}$ is its upper certainty equivalent. A *representor* is a function $\rho : \mathcal{A} \to \mathcal{A} \times \mathcal{A}$. It can be understood as assigning to every act f a place (that of the pair $\rho(f)$) in the stakes order. Two representors shall be of interest:

$$\rho_{MCE}(f) = \begin{cases} (f, \underline{c_f}) & \text{if } (f, \underline{c_f}) \geqq (f, \overline{c_f}) \\ (f, \overline{c_f}) & \text{otherwise} \end{cases}$$

$$\rho_{LCE}(f) = (f, \underline{c_f}).$$

The Maximal Certainty Equivalent representor $\rho_{MCE}(f)$ assigns to an act f the place in the stakes order that is highest out of the stakes involved in the choice between the act and its lower certainty equivalent, and those involved in the choice between it and its upper certainty equivalent. The Lower Certainty Equivalent representor $\rho_{LCE}(f)$ treats an act as if it has the same stakes as those involved in the choice between the act and its lower certainty equivalent.

We are interested in completions satisfying transitivity, as well as the standard Archimedean continuity condition.[17]

Theorem 3. *Let $\preceq$ satisfy A1–A9, and be represented according to (1) by (u, Ξ, D). Then*

(i) *if $\preceq^c$ is a transitive Archimedean completion of $\preceq$, then there exists $\alpha : \mathcal{A} \to [0,1]$ such that, for all $f, g \in \mathcal{A}$, $f \preceq^c g$ iff $V(f) \leqslant V(g)$, where*

$$V(f) = \alpha(f) \min_{p \in D(\rho_{MCE}(f))} \sum_{s \in S} u(f(s)).p(s) + (1-\alpha(f)) \max_{p \in D(\rho_{MCE}(f))} \sum_{s \in S} u(f(s)).p(s) \quad (4)$$

(ii) *if $\preceq^c$ is a transitive Archimedean completion of $\preceq$ satisfying B2^{S-N}, then there exists $\alpha : \mathcal{A} \to [0,1]$ such that, for all $f, g \in \mathcal{A}$, $f \preceq^c g$ iff $V(f) \leqslant V(g)$, where*

$$V(f) = \alpha(f) \min_{p \in \bigcap_{C \in \Xi} C} \sum_{s \in S} u(f(s)).p(s) + (1-\alpha(f)) \max_{p \in \bigcap_{C \in \Xi} C} \sum_{s \in S} u(f(s)).p(s). \quad (5)$$

Part (i) of this result reveals the consequences of demanding that the completion is transitive. It tells us that the completed preferences are represented by a functional similar to the generalised Hurwicz representation (Ghirardato et al., 2004; Cerreia-Vioglio et al., 2011), which involves a mixture of the maximum and minimum expected utility taken over a set of probability measures. However, unlike this representation, the set of probability measures involved depends on the act evaluated. In particular, interpreting $\rho_{MCE}(f)$ as giving the stakes level corresponding to f, the set depends on the stakes involved in the choice of the act. So, representation (4), like representation (2), can be thought of as employing a strategy for preference completion that respects confidence. The difference lies in the domain of the notion of stakes used. Representation (2) uses the stakes involved in the *choice between two acts* to determine the level of confidence appropriate for the evaluation of each of the acts. Representation (4) uses the stakes associated with the *act itself* (via the representor ρ_{MCE}) to fix the appropriate confidence level for its evaluation. Theorem 3 part (i) shows that effectively using a stakes relation on acts, appropriately derived from the notion of stakes over choices, to set the confidence level is the only way to guarantee that the completion is transitive.

Part (ii) shows that this discussion is orthogonal to the issue of the distinction between completion strategies introduced above. Representation (5) uses the set of probability measures encapsulating all of the decision maker's beliefs, even those she holds with little confidence. So, just like B3 and B3^{S-N} (see Theorem 2), the difference between Consistency (B2, which is by definition satisfied by completions; Section 4.1) and its stronger stakes-neutral version, B2^{S-N}, characterises precisely the difference between a strategy that respects confidence and one that goes on hunches.

Note that this result—including the distinction between the two strategies—can be extended to yield a characterisation of completions using the maxmin expected utility rule, rather than the generalised Hurwicz one. As the following proposition shows, it suffices to add the Caution axiom proposed by Gilboa et al. (2010) or its stakes-neutral version, respectively.

Axiom B4 (*Caution*)**.** For all $f \in \mathcal{A}$, $c \in \Delta(X)$, if $f \not\succeq c$, then $c \succeq^c f$.

Axiom B4^{S-N} (*Stakes-Neutral Caution*)**.** For all $f, g \in \mathcal{A}$, if $f' \not\succeq c'$ for all $(f', c') \in \widehat{(f, c)}$, then $c \succeq^c f$.

Proposition 3. *Let $\preceq$ satisfy A1–A9, and be represented according to (1) by (u, Ξ, D). Then*

(i) *if $\preceq^c$ is a transitive Archimedean completion of $\preceq$ satisfying B4, then, for all $f, g \in \mathcal{A}$:*

$$f \preceq^c g \quad \text{iff} \quad in_{p \in D(\rho_{LCE}(f))} \sum_{s \in S} u(f(s)).p(s) \leqslant \min_{p \in D(\rho_{LCE}(g))} \sum_{s \in S} u(g(s)).p(s) \quad (6)$$

(ii) *$\preceq^c$ is a transitive Archimedean completion of $\preceq$ satisfying B2^{S-N} and B4^{S-N} if and only if $\preceq$ is represented according to (3).*

Note that Benchmark on Certainty (B3) entails B4. So part (i) of this proposition implies that whenever B3 is satisfied and $\preceq^c$ is transitive, the completion will conform to a stakes-sensitive version of the maxmin expected utility representation, where the appropriate level of confidence for the evaluation of an act depends on the stakes associated with it (via ρ_{LCE}), independently of the pairwise comparison in which it is considered. Part (ii), which characterises how the distinction between the two strategies plays out in the context of this completion procedure, implies that Stakes-neutral Benchmark on Certainty (B3^{S-N}) is equivalent to the stakes-neutral B2^{S-N} and B4^{S-N}, when the completion satisfies standard conditions.

Remark 2. Some of these results can be sharpened for specific notions of stakes with particular properties. Take, for instance, notions of stakes that always rank choices between an act and a constant act as having low stakes, in the following sense: for every act $f \in \mathcal{A}$ and constant act $c \in \Delta(X)$ that neither strictly dominates nor is strictly dominated by f (i.e. neither $c \succ f(s)$ for all $s \in S$ nor $c \prec f(s)$ for all $s \in S$), $(f, h) \geqq (f, c)$ for every $h \in \mathcal{A}$. (The notions (i) and (ii) in Fig. 2 have this property.) For such notions of stakes, the conditions in Proposition 3 part (i) are both necessary and sufficient for representation (6).

Remark 3. Some notions of stakes over binary choices are naturally built from notions of stakes over acts themselves; in such cases, the representors used above may 'recover' the initial notion of stakes over acts. For example, notion (i) in Fig. 2—stakes as the utility of the least preferred consequence taken over both acts—is naturally obtained from the notion of stakes on acts which looks at the least preferred consequence the act could yield; and both representors above generate precisely this notion of stakes on acts. Moreover, this is the notion effectively used by Hill (2013). Indeed, under notion (i) in Section 2.2, Proposition 3 part (i) (and Remark 2) provides a novel foundation for the confidence-based representation in Hill (2013), in much the same way that Gilboa et al. (2010) can be thought of as providing a novel foundation for the maxmin EU representation of Gilboa and Schmeidler (1989).

17 Namely, for all $f, g, h \in \mathcal{A}$, the sets $\{\alpha \in [0,1] \mid f_\alpha h \preceq g\}$ and $\{\alpha \in [0,1] \mid f_\alpha h \succeq g\}$ are closed in $[0, 1]$.

Finally, let us comment on the case of a decision maker whose confidence ranking contains a singleton set.[18] Such a decision maker is a 'Bayesian with confidence': if forced to give her best estimate for the probability of any event, she could come up with a single value (and these values satisfy the laws of probability), although she may not be very confident in it. Faced with a decision in which she is forced to choose, but where she has little confidence in the relevant beliefs, if such a decision maker goes on hunches, then she acts precisely like a subjective expected utility maximiser (under any of the completion procedures considered above). 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.

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 (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 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 $\mathcal{A}$.[19] 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 $i = 1 \dots n$. Each has preferences $\preceq^i$ over $\mathcal{A}$ (and hence over $\Re^S_+$) represented as in (1) for a stakes relation $\leqq^i$. Each agent thus has a utility function $u^i : \Re_+ \to \Re$, a balanced continuous 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 (1), the economy may contain agents with these sorts of preferences. The aggregate endowment is $e \in \Re^S_{++}$. Finally, an allocation $(x^1, \dots, x^n) \in (\Re^S_+)^n$ 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 3. An allocation $(y^1, \dots, y^n)$ *Pareto dominates* the allocation $(x^1, \dots, x^n)$ if, for each agent i, either $y^i \succ^i x^i$ or $y^i = x^i$.

A feasible allocation $(x^1, \dots, x^n)$ 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 exactly 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.

In Appendix A, we provide a general characterisation result for Pareto optimality, under some technical assumptions. A direct corollary is that, when the aggregate endowment is constant across states, an interior full insurance allocation $(x^1, \dots, x^n)$ is Pareto optimal if and only if $\bigcap_i \bigcap_{z \neq x^i} ri(D^i((x^i, z))) \neq \varnothing$ (see Corollary 2 in Appendix A). It follows that, *grosso modo*, if each agent in the economy who requires more confidence to take decisions with higher stakes were to simply ignore 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.[20] So an economy with agents represented by (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 a set corresponding to the lowest stakes level for choices involving the commodity she is allocated. This highlights some similarities between economies with agents à la Bewley and those with agents represented according to (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 4, Pareto optima correspond to appropriately defined equilibria, this is closely related to the question of how fast the economy can arrive at equilibrium.)[21] There is often a simple, if idealised, fastest way to achieve a Pareto optimum. In particular, whenever a non-Pareto optimal allocation $(x^1, \dots, x^n)$ is dominated by a (feasible) Pareto optimal one $(y^1, \dots, y^n)$, 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 among the agents.) Whenever this is the case, we say that $(y^1, \dots, y^n)$ is *one-step accessible* from $(x^1, \dots, x^n)$. 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), every Pareto dominated allocation is Pareto dominated by a Pareto optimum. In other words, every allocation has a Pareto optimum that is one-step accessible from it. This is not necessarily true for economies containing agents represented by (1), as the following example shows.

Example 2. 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

[18] Hill (2013) calls such confidence rankings 'centred'.

[19] 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 $\Re_+$ (i.e. 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 2, and in particular the generic symbols $f, g, h \dots$ for acts; we use standard vector notation and generic symbols $x, y, z \dots$ for commodities.

[20] 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 $\bigcap_{i \in I} ri(\mathcal{C}_i) = ri \bigcap_{i \in I} \mathcal{C}_i$.

[21] 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 (1).

are represented by (1), where the stakes in the choice between x and y with $x \neq y$ are given by $\max_s |x(s) - y(s)|$.[22] She has the following confidence ranking: $\{\{p \in \Delta(S) | \; 0.5 - \epsilon \leqslant p(s_1) \leqslant 0.5 + \epsilon\} | \; \epsilon \in (0, 0.45]\}$. Note that since each set in the confidence ranking is uniquely identified by an $\epsilon \in (0, 0.45]$, the cautiousness coefficient is entirely specified by a function from pairs 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 all $x \neq y$ and for some $\eta > 0$, where η characterises the agent's attitude to choosing in the absence of confidence (see Section 3.4). Agent 2 is an expected utility decision maker with probability measure assigning 0.5 to both states.

Suppose that there is 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 can be shown that the only Pareto optima are full insurance allocations (see Theorem 4 in Appendix A). 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 they 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 (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[23]

$$p \leqslant \frac{\frac{1}{2}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}} \quad \text{for all } p \in D^1\left(\left(x^1, \frac{1}{2}w\right)\right). \tag{7}$$

Hence, by the definition of D^1, $\frac{1}{2}w \not\succ^1 x^1$ whenever

$$\min\left\{\eta w\left(\delta - \frac{1}{2}\right), 0.45\right\} + 0.5 > \frac{\frac{1}{2}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}. \tag{8}$$

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). By a similar argument, one can show that no Pareto optimal allocation is one-step accessible from (x^1, x^2). Hence the following result.

Proposition 4. *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 (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 smallest 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 (1), she may not be confident enough in her 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 she 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 she 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 had been 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 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.

[22] As noted in footnote 19, although the stakes relation is defined on pairs of contingent commodities, this yields a well-defined stakes relation on pairs of acts. Given the utility function, it is straightforward to check that this stakes relation satisfies the properties assumed in this paper.

[23] Here we consider the case where $\gamma_1 \neq 1$; the case of $\gamma_1 = 1$ can be treated similarly.

Proposition 5. *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 reaching a Pareto optimum, where all agents have sufficient confidence to accept all the trades. 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.

These results may be relevant in situations where there are obvious advantages from reallocation—be it by trade in a market, or by policy implementation by a government—but changes arise slowly, if not at all. For instance, a recognised 'puzzle' is that reforms that most economists agree to be socially beneficial are either not implemented, or deemed unpopular (Williamson, 1994; Rodrik, 1996; Olofsgard, 2003). Many of the most salient examples discussed in the literature follow some shock (e.g. the crises in Latin America of the 1970–80's, the fall of communism in Eastern Europe), after which one might expect the public to have low confidence in their beliefs. The previous propositions suggest that this limited confidence could lead the public to resist big, high-stakes reforms—even ones which they would agree with if they had more confidence or if the stakes were lower. The confidence model thus provides a novel explanation of this puzzle. Moreover, it suggests that small, lower-stakes reforms that 'move' in the direction of the big one may be feasible, though they might not allow the policy maker to arrive at a Pareto optimum. Another consequence of these findings concerns market behaviour after a large shock (such as the 2008 crash). The propositions identify a theoretical friction preventing a market from coming to equilibrium, to be contrasted with practical factors (concerning market structure, for instance) that might slow speed of convergence (or prevent it completely). They suggest that in cases where investor confidence is low (such as after a shock), markets may be slower to come to equilibrium than in other situations. Whilst this appears to be consistent with anecdotal evidence (e.g. on the state of the markets since 2008), we know of no existing empirical work testing this prediction.

6. Related literature

Bewley (1986/2002) 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 Ghirardato et al. (2004) and Gilboa et al. (2010), who take the weak rather than the strict preference relation as primitive.[24] 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 (1) can thus be thought of as 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 (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: firstly, the decision rule which determines preferences on the basis of a set of probability measures and a utility function, and, secondly, the notion of stakes. In this perspective, the current paper can be thought of as complementary to Hill (2013), exploring different parts of the family introduced there. Theorem 1 axiomatises the models in the family that take the unanimity decision rule (as opposed to the maxmin expected utility rule, as in Hill, 2013) and any notion of stakes over binary choices satisfying some basic properties (as opposed to a particular notion of stakes over acts, as in Hill, 2013; see also Section 4.3, Remark 3). Moreover, some of the results in Section 4 (in particular, their part (i)'s) can be thought of as providing foundations for other classes of models belonging to the same family, with different decision rules (maxmin expected utility, generalised Hurwicz) and notions of stakes (over choices, over acts). Hill (2013) discusses the relationship between the proposed family of models incorporating confidence and the existing literature on complete preferences, and in particular ambiguity (such as Klibanoff et al., 2005; Maccheroni et al., 2006; Chateauneuf and Faro, 2009). We refer the reader to that paper for more details on the comparison with models of complete preferences, and restrict our discussion here to the literature on incomplete preferences.

Nau (1992) has proposed a theory of incomplete preferences which is similar to representation (1) in content and motivation. Besides the differences in framework (he uses the de Finetti framework, rather than the Anscombe–Aumann one used here), presentation (he uses confidence-weighted upper and lower conditional probabilities on random variables, rather than the notions of confidence ranking and cautiousness coefficient) and conceptualisation (the distinction between stakes and confidence is not fully brought out; the notion of cautiousness coefficient, and with it the separation of confidence in beliefs from attitudes to choosing in the absence of confidence, is absent), he assumes a particular notion of stakes, whereas we do not. In fact, Nau's model is the special case of the theory presented here with (iv) in Fig. 2 as stakes. His particular notion of stakes imposes properties on the representation, such as convexity (Nau, 1992, p1741), which do not apply for all notions; whilst they play a central role in his axiomatisation, they are thus absent from ours. On the other hand, our result brings out some general aspects of the representation—for example, the fact that it respects Pure Independence, and hence, as argued in Sections 1 and 3.2, the essence of the standard independence axiom—that lie beyond the focus of Nau's presentation and motivation. Finally, Nau does not discuss the implications of confidence for preference completion, which was treated in Section 4.

Recently, Faro (2015) has proposed an extension of Bewley's representation incorporating an ambiguity index—a real-valued function on the space of probability measures—in a way inspired by the variational preferences model of Maccheroni et al. (2006). His representation, like Bewley's and the one studied here, involves a universal quantification over probability measures; Cerreia-Vioglio et al. (2015) axiomatise the 'justifiable' version of his model, involving an existential quantification instead. Lehrer and Teper (2011, Theorem 2) axiomatise a representation involving sets of sets of probability measures, where an act is preferred to another if it has a higher expected utility for all probability

24 The representation in Ghirardato et al. (2004) and Gilboa et al. (2010) differs from representation (1) by replacing $D((f, g))$ with a fixed set of probability measures; the representation in Bewley (1986/2002) differs moreover in replacing the weak preferences and orders by strict ones.

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 modelled, and the behavioural properties. Conceptually, Faro (2015) suggests an interpretation of his ambiguity index in terms of confidence (in experts or opinions). However, his model does not support a distinction between beliefs and tastes (see his Section 4.3), whereas the difference between the confidence ranking and the cautiousness coefficient in the model proposed here can be considered to correspond precisely to such a distinction (Sections 2.3 and 3.4). As noted in Section 2.3, this separation is essential to the concept of confidence in beliefs that is at issue in this paper. (See also Hill, 2013 for a discussion of a similar point in relation to ambiguity models.) On the behavioural front, these models employ more severe weakenings of transitivity than used here, and indeed allow preference cycles whilst representation (1) does not (Section 3.2); moreover, Faro (2015) employs a more severe weakening of independence. Faro (2015) 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 4.

Minardi and Savochkin (2015) 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 3.4, with the notable difference that whilst Minardi and Savochkin (2015) assume this relation as a primitive, here it is defined from (ordinary) preferences over acts (Definition 2, Section 3.4). 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 (2015) and Lehrer and Teper (2011) than to the one used here.

Seidenfeld et al. (1995), Nau (2006), Ok et al. (2012) and Galaabaatar and Karni (2013) have explored extensions of Bewley's representation involving sets of probabilities and sets of utilities. The behavioural points made in the Introduction, 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 (2012) 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. Hill (2016) develops a model of choice in the presence of a costly deferral option which contains a natural extension of representation (1) to general, as opposed to two-element menus, where stakes are assigned to menus and the only acts chosen from a menu have highest expected utility according to the appropriate set of priors (see his Section 3). Moreover, as opposed to the approach adopted here (Section 2.2), the (analogue of the) stakes relation is endogenously derived (see his Remark 1).

The discussion of the completion of incomplete preferences is technically related to Gilboa et al. (2010), Kopylov (2009), Nehring (2009) and Danan et al. (forthcoming), who provide results relating 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. 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 4. Putting aside this point, and notwithstanding some technical differences, our Proposition 3 is closely related to Theorem 3 in Gilboa et al. (2010), whereas our Theorem 3 can be thought of as a version of Proposition 2 in Danan et al. (forthcoming). By contrast, Theorem 2 involves a new sort of 'connecting' axiom (Benchmark on Certainty) and provides, in the case of a degenerate confidence ranking (containing a single set of probability measures), a new axiomatisation of the representation obtained in Gilboa et al. (2010)'s Theorems 3 and 4.

Finally, the interpretation of indeterminacy of preference in terms of sticking to a status quo option used in Section 5 has been considered by Bewley (1986/2002), 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.

7. 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, suggests that the confidence ranking captures the decision maker's confidence in beliefs, and the cautiousness coefficient her 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 and behaviourally 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 mobilises 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 axiomatic characterisations of several completion procedures using the two strategies were proposed.

Finally, possible consequences of the model in a market setting were considered, where indeterminacy of preferences translates into refusal to trade. In particular, 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.

Appendix A. Characterisation of Pareto optima under representation (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 5. This is in fact not straightforwardly possible in general, for two reasons. Firstly, the weakening of transitivity in representation (1) implies that, for certain notions of stakes, preferences represented by (1) may not be convex.[25] Secondly, since the weak preference order is taken as primitive, it does not follow from the axioms in Section 3.2 that the strict preference order is continuous.[26]

To deal with these issues, we assume the following properties of the notion of stakes and of the preferences.

Monotone decreasing For all $\preceq_{\Delta(X)} \in \mathcal{P}_{\Delta(X)}, f, g \in \mathcal{A}$ and $\alpha, \beta \in [0,1]$, if $\alpha \leqslant \beta$, then $(f, g_\alpha f) \leqq_{\preceq_{\Delta(X)}} (f, g_\beta f)$.

Full support For all $s \in S$, there exists $c_s \in \Re_{++}$ such that, for all $h \in \mathcal{A}$ and $\alpha \in (0,1]$, $(1_s)_\alpha h \succ (c_s)_\alpha h$.[27]

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 Section 2.2.

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.

Under these assumptions, we have a characterisation of Pareto optima. For a contingent commodity $x \in \Re^S_+$ and an agent i, let

$$\Pi^i(x) = \left\{ \left(\frac{p(s_1)u^{i\prime}(x_{s_1})}{\sum_{t \in S} p(t)u^{i\prime}(x_t)}, \ldots, \frac{p(s_{|S|})u^{i\prime}(x_{s_{|S|}})}{\sum_{t \in S} p(t)u^{i\prime}(x_t)} \right) \middle| \; p \in \bigcap_{z \neq x} ri(D^i((x,z))) \right\}.$$

[25] 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 $g, h \succ f \not\succeq g_\alpha h$ are compatible with representation (1).

[26] 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 (1986/2002); see Section 6.

[27] 1_s is the characteristic function for s: $1_s(s') = 1$ for $s' = s$ and $1_s(s') = 0$ for $s' \neq s$. As specified in Section 2.1, c_s is the constant act taking the degenerate lottery yielding c_s for sure in all states.

Theorem 4. *Suppose that, for each $1 \leqslant i \leqslant n$, the notion of stakes $\leqq^i_\bullet$ is monotone decreasing and the preference relation $\preceq^i$ satisfies full support. 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 2. *Suppose that the aggregate endowment is constant across states and that, for each $1 \leqslant i \leqslant n$, the notion of stakes $\leqq^i_\bullet$ is monotone decreasing and the preference relation $\preceq^i$ satisfies full support.*

(i) *An interior full insurance allocation $(x^1, \ldots, x^n)$ is Pareto optimal iff $\bigcap_i \bigcap_{z \neq x^i} ri(D^i((x^i, z))) \neq \varnothing$.*
(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 $\bigcap_i ri(\mathcal{C}^i) \neq \varnothing$. 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 (1). It is in the style of existing results, such as Billot et al. (2000), Rigotti and Shannon (2005) and 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 (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 the representation axiomatised by Bewley (1986/2002) (see Section 6), and take a stricter notion of Pareto optimality.

Appendix B. 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.

B.1. Proof of Theorem 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–Morgenstern 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 B.5) that, for each non-minimal stakes level r, $\preceq_r$ is a non-degenerate, 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 B.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 B.7, the $\mathcal{C}_r$ form a nested family of sets, and we thus have a confidence ranking. By Lemmas B.8 and B.9, this confidence ranking is continuous; by Lemma B.11, it is balanced. 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 (A3) is equivalent to the standard transitivity axiom, and A4 and 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 B.1. *There exists a non-constant utility function u representing the restriction of $\preceq$ to the constant acts. Moreover, $u(\Delta(X)) = \Re$.*

Proof. By A1, A2, A4 and A7, the restriction of $\preceq$ to constant acts is non-degenerate complete, reflexive and satisfies independence and continuity. We now show that it is transitive. For any $c, d, e \in \Delta(X)$, suppose that $c \preceq d$ and $d \preceq e$. If (c,e) is $\leqq$-minimal, then (A3) immediately implies that $c \preceq e$. Now suppose that (c,e) is not $\leqq$-minimal. If $c \sim d \sim e$, then $c \sim e$ by (A3). If $c \sim d$ and $d \nsim e$, then by the extensionality of $\leqq$, $(c,e) \equiv (d,e)$, so (A3) implies that $c \preceq e$. The case where $d \sim e$ and $c \nsim d$ is treated similarly. Consider finally the case where $c \nsim d$ and $d \nsim e$. By the richness of $\leqq$, there exist $(c',d') \in \widehat{(c,d)}$ and $(d'',e'') \in \widehat{(d,e)}$ such that $(c,e) \leqq (c',d')$ and $(c,e) \leqslant (d'',e'')$. By A1 and A4, $c'(s) \preceq d'(s)$ and $d''(s) \preceq e''(s)$ for all $s \in S$, from which it follows by (A5) that $c' \preceq d'$ and $d'' \preceq e''$. Hence, by (A3), $c \preceq e$, as required. The existence of u follows from the von Neumann–Morgenstern theorem. The unboundedness of u is a straightforward consequence of A9. □

There is thus a many-to-one mapping between acts in $\mathcal{A}$ and elements of B, 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 by $\preceq$ under this mapping, and $\leqq$ to denote the order generated on $B \times B \setminus \{(a,a) \mid a \in B\}$ by $\leqq$. ($\preceq$ and $\leqq$ are well-defined on B by A3, A5 and the extensionality of $\leqq$) Similarly, we use $\widehat{(\bullet,\bullet)}$ to denote the mapping on $B \times B$ generated by $\widehat{(\bullet,\bullet)}$: explicitly, for $(a,b) \in B \times B$, $\widehat{(a,b)} = \{(a',b') \in B \times B \mid \exists \alpha > 0,\ l \in B\ s.t.\ a' = \alpha a + (1-\alpha)l \text{ and } b' = \alpha b + (1-\alpha)l\}$.

Lemma B.2. *For every $a, \bar{a}, b, \bar{b} \in B$ with $a \neq b$ and $\bar{a} \neq \bar{b}$, there exists $(a',b') \in \widehat{(a,b)}$ such that $(a',b') \equiv (\bar{a},\bar{b})$.*

Proof. If $(a,b) \equiv (\bar{a},\bar{b})$, then there is nothing to show. Suppose without loss of generality that $(a,b) < (\bar{a},\bar{b})$; the other case is treated similarly. By richness of $\leqq$ and the definition of $\widehat{(\bullet,\bullet)}$, there exist $\beta > 0$ and $l \in \mathcal{A}$ such that $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \geqq (\bar{a},\bar{b})$. If $(\beta a + (1-\beta)l, \beta b + (1-\beta)l) \equiv (\bar{a},\bar{b})$, then the result has been established; if not, then by continuity of $\leqq$, there exists $\alpha > 0$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \equiv (\bar{a},\bar{b})$. Since $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in \widehat{(a,b)}$, this yields the required result. □

Let $\mathcal{S}$ be the set of equivalence classes of $\leqq$. As standard, $\leqq$ on $B \times B \setminus \{(a,a) \mid a \in B\}$ 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 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 element; if it exists, we denote the minimal element by $\underline{\mathcal{S}}$. $r \in \mathcal{S}$ is full if, for every $a,b \in B$ with $a \neq b$, there exists $(a',b') \in \widehat{(a,b)}$ such that $(a',b') \in r$. It follows from Lemma B.2 that every element in $\mathcal{S}$ is full. Let $\mathcal{S}^+$ be the set of non-minimal elements. For each $r \in \mathcal{S}^+$, let $\preceq_r$ be the reflexive binary relation on B such that, for all $a,b \in B$ with $a \neq b$, $a \preceq_r b$ iff there exists $(a',b') \in \widehat{(a,b)}$ such that $(a',b') \in r$ and $a' \preceq b'$. The following lemma implies that, for $a,b \in B$ with $a \neq b$ and every $r \in \mathcal{S}^+$, $a \preceq_r b$ iff $a' \preceq b'$ for every $(a',b') \in \widehat{(a,b)}$ such that $(a',b') \in r$.

Lemma B.3. *For every $a, b, l, m \in B$ and $\alpha, \beta > 0$ with $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l), (\beta a + (1-\beta)m, \beta b + (1-\beta)m) \in r \in \mathcal{S}^+$, $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$ iff $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$.*

Proof. Since $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$, $a \neq b$. Without loss of generality, suppose that $\beta \leqslant \alpha$. We first establish the result for 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 $\frac{\alpha\beta-\beta}{\alpha-\beta}l + \frac{\alpha-\alpha\beta}{\alpha-\beta}m \in B$; 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 (A6), $f \asymp g$ iff $f_{\frac{\beta}{\alpha}}h \asymp g_{\frac{\beta}{\alpha}}h$. Hence, by 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 $l \neq m$. Suppose that $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$; we show that $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$. Since r is non-minimal, there exists $s \in \mathcal{S}^+$ with $s < r$. Since s is full, there exists $n \in B$ and $\gamma > 0$, $\gamma \neq 1$, such that $(\gamma(\alpha a + (1-\alpha)l) + (1-\gamma)n, \gamma(\alpha b + (1-\alpha)l) + (1-\gamma)n) \in s$. Let $X = (\frac{\gamma\alpha\beta-\beta}{\gamma\alpha-\beta}(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n) + \frac{\gamma\alpha-\gamma\alpha\beta}{\gamma\alpha-\beta}m)$ and consider $\delta(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X$ for $\delta > 0$. Note that this equals $\gamma(\alpha a + (1-\alpha)l) + (1-\gamma)n$ when $\delta = 1$ and $\beta a + (1-\beta)m$ when $\delta = \frac{\beta}{\gamma\alpha}$; moreover, it can be rewritten as $\delta\gamma(\alpha a + (1-\alpha l)) + (1-\delta\gamma)(\frac{\delta-\delta\gamma}{1-\delta\gamma}n + \frac{1-\delta}{1-\delta\gamma}X)$. Hence, by A6 and A4 and since $\alpha a + (1-\alpha)l \preceq \alpha b + (1-\alpha)l$, for every $\delta \neq \frac{1}{\gamma} = \frac{\beta}{\alpha\gamma}$ such that $(\delta(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X, \delta(\alpha\gamma b + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X) \leqq (\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l)$, we have that $\delta(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X \preceq \delta(\alpha\gamma b + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X$. By the continuity of $\leqq$, there exists a limit $\bar{\delta}$ with $(\bar{\delta}(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\bar{\delta})X, \bar{\delta}(\alpha\gamma b + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\bar{\delta})X) \equiv (\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l)$ and $(\delta(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\bar{\delta})X, \bar{\delta}(\alpha\gamma b + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\delta)X) \leqq (\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l)$ for all $\delta \in (\frac{\beta}{\alpha\gamma}, 1)$. If $\frac{\beta}{\alpha\gamma}$ is such a limit $\bar{\delta}$, then it follows by A7 that $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, as required. If not, then take any such limit $\bar{\delta}$: by the previous observation, $\bar{\delta}(\alpha\gamma a + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\bar{\delta})X \preceq \bar{\delta}(\alpha\gamma b + (1-\gamma\alpha)(\frac{\gamma-\gamma\alpha}{1-\gamma\alpha}l + \frac{1-\gamma}{1-\gamma\alpha}n)) + (1-\bar{\delta})X$, with $\bar{\delta} \neq \frac{\beta}{\alpha\gamma}$. It follows by A6 and A4 that $\beta a + (1-\beta)m \preceq \beta b + (1-\beta)m$, as required. The converse is established by the same argument. □

We now establish some properties of the relations $\preceq_r$.

Lemma B.4. *For all $r, s \in \mathcal{S}^+$ with $r \geqslant s$, $\preceq_r \subseteq \preceq_s$.*

Proof. If $s = r$, there is nothing to show, so suppose not. 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 B.3 that $a \preceq b$. Let $\beta > 0$ and $m \in B$, be such that $(\beta a + (1-\beta)m, \beta b + (1-\beta)m) \in s$ (such β and m exist since s is full). A6 and 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 is

- *non-degenerate* if there exists $a, b \in B$ such that $a \preceq b$ but not $a \succeq b$.
- *monotonic* if, for all $a, b, \in B$, if $a \leqslant b$ then $a \preceq b$.
- *affine* if, for all $a, b, c \in B$ and $\alpha \in (0, 1)$, $a \preceq b$ iff $\alpha a + (1-\alpha)c \preceq \alpha b + (1-\alpha)c$.
- *Archimedean* if, for all $a, b, c \in B$, 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 B.5. *For every $r \in \mathcal{S}^+$, $\preceq_r$ is a non-degenerate, monotonic, affine, Archimedean pre-order.*

Proof. *Non-degeneracy*. By A2, $\preceq$ is non-degenerate; by A5 and A1, it follows that the restriction of $\preceq$ to $\Delta(X)$ is non-degenerate. But $\preceq_r$ coincides with $\preceq$ on $\Delta(X)$, so it is non-degenerate.

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 $l \in B$ and $\alpha > 0$ such that $(\alpha a + (1-\alpha)l, \alpha b + (1-\alpha)l) \in r$. By monotonicity (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 is full, there exists $\beta > 0$ and $l \in B$ 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 is full, there exists $\gamma > 0$ and $m \in B$ 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 $\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m \in B$, and similarly for b. We now distinguish three cases.

If $\gamma\alpha < 1$, 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 A6 and 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.

If $\gamma\alpha > 1$, then the same argument can be applied, with $f, g, h \in \mathcal{A}$ such that $\alpha\gamma(\beta a + (1-\beta)l) + (1-\alpha\gamma)(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m) = u \circ f$, $\alpha\gamma(\beta b + (1-\beta)l) + (1-\alpha\gamma)(\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m) = u \circ g$ and $\frac{\gamma-\alpha\gamma}{1-\alpha\gamma}(\beta c + (1-\beta)l) + \frac{1-\gamma}{1-\alpha\gamma}m = u \circ h$, and using $f_{\frac{1}{\alpha\gamma}}h$ and $g_{\frac{1}{\alpha\gamma}}h$.

The final case is when there exists no $\gamma \neq \frac{1}{\alpha}$ satisfying the conditions stated above. This case is treated analogously to the $\alpha = \beta$ case in the proof of Lemma B.3.

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 is full, there exists $l \in B$ and $\alpha > 0$ such that $(\alpha a + (1-\alpha)l, \alpha c + (1-\alpha)l) \in r$. Moreover, there exists $m, n \in B$ and $\beta, \gamma > 0$ 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 $b \preceq_r c$, by Lemma B.3, $\beta(\alpha a + (1-\alpha)l) + (1-\beta)m \preceq \beta(\alpha b + (1-\alpha)l) + (1-\beta)m$ and $\gamma(\alpha b + (1-\alpha)l) + (1-\gamma)m \preceq \gamma(\alpha c + (1-\alpha)l) + (1-\gamma)m$. Hence, by A3, $\alpha a + (1-\alpha)l \preceq \alpha c + (1-\alpha)l$, and so $a \preceq_r c$, as required.

Archimedean. 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 versions 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 B.4, $\beta a + (1-\beta)b \succeq c$ for all $\beta \in I$, whence $\bar{\alpha}a + (1-\bar{\alpha})b \succeq c$ by 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 B.2, there exists $l \in B$ and $\bar{\delta} > 0$ such that $(\bar{\delta}(\bar{\alpha}a + (1-\bar{\alpha})b) + (1-\bar{\delta})l, \bar{\delta}c + (1-\bar{\delta})l) < (\bar{\alpha}a + (1-\bar{\alpha})b, c)$. Suppose that $\bar{\delta} < 1$; the other case is treated similarly. 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 the 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 B.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 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 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 B.6. *For each $r \in \mathcal{S}^+$, there exists a unique closed convex set of probabilities $\mathcal{C}_r$ such that, for all $a, b \in B$, $a \preceq_r b$ iff*

$$\sum_{s \in S} a(s)p(s) \leqslant \sum_{s \in S} b(s)p(s) \quad \text{for all } p \in \mathcal{C}_r. \tag{B.1}$$

Proof. This follows from Lemma B.5, by Gilboa et al. (2010, Corollary 1),[28] which establishes such a representation for non-degenerate, monotonic, affine, Archimedean pre-orders. □

Lemma B.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 B.4 and Ghirardato et al. (2004, Proposition A.1). □

Lemma B.8. *For all $r \in \mathcal{S}^+$, $\mathcal{C}_r = \overline{\bigcup_{r' < r} \mathcal{C}_{r'}}$.*

Proof. By Lemma 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) $p \in \mathcal{C}_r \setminus \overline{\bigcup_{r' < r} \mathcal{C}_{r'}}$. 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 $q \in ba(S)$. Since r is full, there exists $\delta > 0$ and $m \in B$ such that $(\delta a + (1-\delta)m, \delta\alpha^* + (1-\delta)m) \in r$. By Lemma B.2, there exists $l \in B$ and $\beta > 0$ such that $(\beta(\delta a + (1-\delta)m) + (1-\beta)l, \beta(\delta\alpha^* +$

[28] See Ghirardato et al. (2004, Proposition A.2) for a related result.

$(1-\delta)m)+(1-\beta)l) < (\delta a+(1-\delta)m, \delta\alpha^*+(1-\delta)m)$. Consider the case in which $\beta < 1$; the other case is treated similarly. 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 (A7). Hence $\mathcal{C}_r = \overline{\bigcup_{r'<r} \mathcal{C}_{r'}}$. □

Lemma B.9. *For all non-maximal $r \in \mathcal{S}^+$, $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$.*

Proof. By Lemma 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) $p \in \bigcap_{r'>r} \mathcal{C}_{r'} \setminus \mathcal{C}_r$. 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 $q \in ba(S)$. Since r is full, there exists $\delta > 0$ and $m \in B$ such that $(\delta a+(1-\delta)m, \delta\alpha^*+(1-\delta)m) \in r$. Take any $x \in \Re$ 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 $s \in S$, for all $(f', (g_\beta h)') \in \widehat{(f, g_\beta h)}$ with $(f', (g_\beta h)') > (f, g)$, and $f' \nsucceq (g_\beta h)'$. Since (f, g) is not $\leqq$-maximal, this contradicts A8; hence $\mathcal{C}_r = \bigcap_{r'>r} \mathcal{C}_{r'}$. □

Lemma B.10. *Let $\preceq_{\bigcap \mathcal{S}}$ be the relation on B generated by (B.1) with the set of probability measures $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r$. If there exists a minimal element of $\mathcal{S}$, $\underline{\mathcal{S}}$, then $\preceq|_{\underline{\mathcal{S}}} = \preceq_{\bigcap \mathcal{S}}|_{\underline{\mathcal{S}}}$.*

Proof. Let $a, b \in B$ 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 (B.1) and Lemma B.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 B.2, there exists $\alpha > 0$ and $l \in B$ such that $(\alpha(\beta a + (1-\beta)x^*) + (1-\alpha)l, \alpha b + (1-\alpha)l) > (a, b)$ and $\alpha(\beta a+(1-\beta)x^*)+(1-\alpha)l \preceq \alpha b+(1-\alpha)l$. Hence, by A8, $a \preceq b$, as required. Now suppose that $a \preceq b$. By 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 B.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. □

Lemma B.11. *For all $r, s \in \mathcal{S}^+$, if $\mathcal{C}_r \subset \mathcal{C}_s$, then $\mathcal{C}_r \cap ri(\mathcal{C}_s) \neq \varnothing$. Similarly, if there exists a minimal element of $\mathcal{S}$, then for all $s \in \mathcal{S}^+$, $\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r \cap ri(\mathcal{C}_s) \neq \varnothing$.*

Proof. We consider only the case of $r, s \in \mathcal{S}^+$; the other case is treated similarly, noting that if $\underline{\mathcal{S}}$ exists, then it is full. Suppose that the condition does not hold, so there exist $r, s \in \mathcal{S}^+$ with $\mathcal{C}_r \subset \mathcal{C}_s$ and $\mathcal{C}_r \cap ri(\mathcal{C}_s) = \varnothing$. Since $\mathcal{C}_r$ is convex and in the relative boundary of $\mathcal{C}_s$, it follows from a supporting hyperplane theorem that there is a linear functional ϕ on $ba(S)$ and $\alpha \in \Re$ and such that $\phi(q) = \alpha$ for all $q \in \mathcal{C}_r$ and $\phi(q) \geqslant \alpha$ for all $q \in \mathcal{C}_s$ with strict inequality for some $q \in \mathcal{C}_s$. 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 $q \in ba(S)$. Since s is full, there exists $\gamma > 0$ and $l \in B$ such that $(\gamma a + (1-\gamma)l, \gamma\alpha^* + (1-\gamma)l) \in s$. Since r is full, there exists $\delta > 0$, $m \in B$ such that $(\delta(\gamma a + (1-\gamma)l) + (1-\delta)m, \delta(\gamma\alpha^* + (1-\gamma)l) + (1-\delta)m) \in r$. Consider the case in which $\delta < 1$; the other case is treated similarly. Let $f, g, h \in \mathcal{A}$ be such that $u \circ f = \gamma a + (1-\gamma)l$, $u \circ g = \gamma\alpha^* + (1-\gamma)l$, $u \circ h = m$. By construction, $f \succ g$, whilst $f_\delta h \sim g_\delta h$, contradicting A4. So there exist no such r, s, as required. □

Conclusion of the proof of Theorem 1. Define

$$\Xi = \begin{cases} \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} & \text{if } \mathcal{S} = \mathcal{S}^+ \\ \{\mathcal{C}_r \mid r \in \mathcal{S}^+\} \cup \Big\{\bigcap_{r\in\mathcal{S}^+} \mathcal{C}_r\Big\} & \text{if } \mathcal{S} = \mathcal{S}^+ \cup \{\underline{S}\} \end{cases}$$

where the $\mathcal{C}_r$ are as specified in Lemma B.6. It follows from Lemma 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 B.6), Ξ is a confidence ranking. By Lemmas B.8 and B.9, Ξ is continuous; by Lemma B.11, it is balanced. D is defined as follows: for all $(f, g) \in \mathcal{A} \times \mathcal{A}$, if $[(f, g)] \in \mathcal{S}^+$, then $D((f, g)) = \mathcal{C}_{[(f,g)]}$; if $(f, g) \in \underline{\mathcal{S}}$, then $D((f, g)) = \bigcap_{s\in\mathcal{S}^+} \mathcal{C}_s$; and if $f(s) \sim g(s)$ for all $s \in S$, then $D((f, g)) = \mathcal{C}$, for some arbitrary $\mathcal{C} \in \Xi$. Order preservation and surjectivity of D are immediate from the definition and Lemma B.7. By construction and Lemma B.10, u, Ξ, D represent $\preceq$ according to (1).

The direction from (ii) to (i) is generally straightforward. The only interesting case is continuity (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 $(\alpha_{i_n}, \beta_{i_n})$, then the result follows from the fact that D is order-preserving and the continuity of the unanimity rule. Now consider the case in which there exists a no such subsequence. 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 (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 $(f, g) \in (\mathcal{A} \times \mathcal{A})^{nt}$ with non-minimal stakes such that $D_1((f, g)) \neq D_2((f, g))$. Suppose, without loss of generality, that $p \in D_1((f, g)) \setminus D_2((f, g))$. 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 $q \in ba(S)$. By the richness of $\leqq$, there exists $\gamma > 0$ and $l \in B$ such that $(\gamma a + (1-\gamma)l, \gamma\alpha^* + (1-\gamma)l) \equiv (u \circ f, u \circ g)$. Taking $h, h' \in \mathcal{A}$ such that $u \circ h = \gamma a + (1-\gamma)l$ and $u \circ h' = \gamma\alpha^* + (1-\gamma)l$, we have that $\sum_{s\in S} u(h(s))p(s) \geqslant \sum_{s\in S} u(h')p(s)$ for all p s.t. $p \in D_2((h, h'))$, whereas this is not the case for all p s.t. $p \in D_1((h, h'))$, 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. □

B.2. Proofs of results in Sections 3.4 and 4

Proof of Proposition 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 the stakes relation and the mapping from $\mathcal{A}$ to B used in the proof of Theorem 1 can be taken to be the same for the two agents; we use the notation employed in that proof. By (i), for every $r \in \mathcal{S}^+$, $\preceq^1_r \subseteq \preceq^2_r$, and so, by Ghirardato et al. (2004, Proposition A.1), $\mathcal{C}^2_r \subseteq \mathcal{C}^1_r$. It follows that $\Xi_2 \sqsubseteq \Xi_1$ and $D_2((f,g)) \subseteq D_1((f,g))$ for all $(f,g) \in (\mathcal{A}\times\mathcal{A})^{nt}$. □

Proof of Proposition 2. The 'if' direction is straightforward. The 'only if' direction is a simple corollary of the proof of Theorem 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 1 are the same, and so the confidence rankings are the same. □

Proof of Theorem 2. First consider part (i). Showing the necessity of the axioms is straightforward; we now show sufficiency. Consider $f, g \in \mathcal{A}$. If $f(s) \sim g(s)$ for all $s \in S$ then $\min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s) = \min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s)$. In this case, B3 implies that neither $f \not\succ^c g$ nor $g \not\succ^c f$, so, by C1, $f \sim^c g$, as required by the desired representation. Suppose henceforth that this is not the case, i.e. that $f(s) \nsim g(s)$ for some $s \in S$. Consider firstly the case where $\min_{D((f,g))} \sum_{s\in S} u(g(s)).p(s) \geqslant \min_{D((f,g))} \sum_{s\in S} u(f(s)).p(s)$. Consider any $c \in \Delta(X)$ such that there exists $(f',c') \in \widehat{(f,c)}$ with $(f',c') \equiv (f,g)$ and $f' \succeq c'$; it thus follows that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(c)$ for all $p \in D((f,g))$. By the richness and continuity of $\leqq$, for each $d \in \Delta(X)$ with $d \prec c$, there exist $(g',d') \in \widehat{(g,d)}$ such that $(g',d') \equiv (f,g)$; moreover, for any such d and (g',d'), $g' \succeq d'$. It follows by B3 that $g \not\prec^c f$, and thus, by C1, $g \succeq^c f$, as required.

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 (1) that for all $(g',d') \in \widehat{(g,d)}$ such that $(g',d') \equiv (f,g)$, $g' \not\preceq d'$. However, we have $u(c) \leqslant \sum_{s\in S} u(f(s)).p(s)$ for all $p \in D((f,g))$. By the richness and continuity of $\leqq$, there exists $(f',c') \in \widehat{(f,c)}$ such that $(f',c') \equiv (f,g)$; by representation (1), $f' \succeq c'$. It follows from B3 that $g \prec^c f$, as required. Hence representation (2) holds.

Now consider part (ii). Showing the necessity of the axioms is straightforward; we show sufficiency. Consider $f, g \in \mathcal{A}$. If $f(s) \sim g(s)$ for all $s \in S$ then $\min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(g(s)).p(s) = \min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$. In this case, B3^{S-N} implies that neither $f \not\succ^c g$ nor $g \not\succ^c f$, so, by C1, $f \sim^c g$, as required by the desired representation. Suppose henceforth that this is not the case, i.e. that $f(s) \nsim g(s)$ for some $s \in S$. Consider firstly the case where $\min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(g(s)).p(s) \geqslant \min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$. Consider any $c \in \Delta(X)$ such that there exists $(f',c') \in \widehat{(f,c)}$ with $f' \succeq c'$; it thus follows that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(c)$ for all $p \in \bigcap_{C\in\Xi} C$. Consider any $d \prec c$. By the continuity of representation (1), the continuity of Ξ and the surjectivity of D, there exists $\mathcal{C} \in \Xi$ such that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(d)$ for all $p \in \mathcal{C}$. Moreover, by the richness of $\leqq$, there exists $(g',d') \in \widehat{(g,d)}$ such that $D((g',d')) \subseteq \mathcal{C}$, so $g' \succeq d'$. Since this holds for all $d \prec c$, it follows by B3^{S-N} that $g \not\prec^c f$, and thus, by C1, $g \succeq^c f$, as required.

Now suppose that $\min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(g(s)).p(s) < \min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$, so there exist $c, d \in \Delta(X)$ with $\min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(g(s)).p(s) < u(d) < u(c) < \min_{\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$. 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 $p \in \bigcap_{C\in\Xi} C$. It thus follows from representation (1), the continuity of Ξ and the surjectivity of D that there exists $(g',d') \in \widehat{(g,d)}$ such that $g' \not\succeq d'$. However, we have $u(c) < \sum_{s\in S} u(f(s)).p(s)$ for all $p \in \bigcap_{C\in\Xi} C$; so, by the continuity representation (1), the continuity of Ξ and the surjectivity of D, there exists $\mathcal{C} \in \Xi$ such that $\sum_{s\in S} u(g(s)).p(s) \geqslant u(c)$ for all $p \in \mathcal{C}$. By the richness and continuity of $\leqq$, there exists $(f',c') \in \widehat{(f,c)}$ such that $D((f',c')) \subseteq \mathcal{C}$, so $f' \succeq c'$. It follows from B3^{S-N} that $g \prec^c f$, as required. Hence representation (3) holds. □

Proof of Theorem 3. First consider part (i). By B1, $\preceq_{\Delta(X)} = \preceq^c_{\Delta(X)}$, so u represents the restriction of $\preceq^c_{\Delta(X)}$ to constant acts. Consider $f \in \mathcal{A}$. Note that $\underline{c_f}$ and $\overline{c_f}$ exist and are uniquely defined up to $\sim$, by representation (1). Suppose without loss of generality that $(f, \underline{c_f}) \geqq (f, \overline{c_f})$, so $\rho_{MCE}(f) = (f, \underline{c_f})$; the other case follows similarly. By the definition, and the continuity of (1), $\underline{c_f} \preceq f \preceq \overline{c_f}$. By representation (1), it follows that $u(\underline{c_f}) \leqslant \sum_{s\in S} u(f(s)).p(s)$ for all $p \in D(\rho_{MCE}(f))$, so $u(\underline{c_f}) \leqslant \min_{p\in D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$. Moreover, if $u(\underline{c_f}) < \min_{p\in D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$, then, by the continuity of Ξ, the surjectivity of D and representation (1), for any $d \in \Delta(X)$ with $d \succ \underline{c_f}$, there exists $\alpha \in (0,1)$ such that $u(\underline{c_f}_\alpha d) \leqslant \min_{p\in D((f,\underline{c_f}_\alpha d))} \sum_{s\in S} u(f(s)).p(s)$, so $\underline{c_f} \prec \underline{c_f}_\alpha d \preceq f$, contradicting the definition of $\underline{c_f}$. So $u(\underline{c_f}) = \min_{p\in D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$. By a similar argument, $u(\overline{c_f}) = \max_{D(f,\overline{c_f})} \sum_{s\in S} u(f(s)).p(s)$. By the order-preserving property of D, it follows that $u(\overline{c_f}) \leqslant \max_{D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$. Take any $d \in \Delta(X)$ such that $u(d) = \max_{D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$ (such a d exists by representation (1)). By construction, $d \succeq \overline{c_f}$, so by (1), $d \succeq f$. It follows from B2 that $\underline{c_f} \preceq^c f \preceq^c d$.

By Archimedean continuity and a standard argument, there exists a unique $\alpha(f) \in [0,1]$ such that $f \sim^c \underline{c_f}_{\alpha(f)} d$. Let $V(f) = u(\underline{c_f}_{\alpha(f)} d)$. By C1 and A3, V represents $\preceq^c$. Since u is affine, $V(f) = \alpha(f)u(\underline{c_f}) + (1-\alpha(f))u(d) = \alpha(f) \min_{D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s) + (1-\alpha(f)) \max_{D(\rho_{MCE}(f))} \sum_{s\in S} u(f(s)).p(s)$, as required.

As for part (ii), by B1, $\preceq_{\Delta(X)} = \preceq^c_{\Delta(X)}$, so u represents the restriction of $\preceq^c_{\Delta(X)}$ to constant acts. Consider $f \in \mathcal{A}$. Let $\underline{c_f}'$ be the $\preceq$-minimal element of $\Delta(X)$ such that, for all $c \in \Delta(X)$ for which there exists $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $f_\alpha h \succeq c_\alpha h$, $\underline{c_f}' \succeq c$ and let $\overline{c_f}'$ be the $\preceq$-maximal element of $\Delta(X)$ such that, for all $c \in \Delta(X)$ for which there exists $\alpha \in (0,1]$ and $h \in \mathcal{A}$ such that $f_\alpha h \preceq c_\alpha h$, $\overline{c_f}' \preceq c$. By B2^{S-N} and reasoning similar to that used in the proof of part (i), $\underline{c_f}' \preceq^c f \preceq^c \overline{c_f}'$, $u(\underline{c_f}') = \min_{p\in\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$ and $u(\overline{c_f}') = \max_{p\in\bigcap_{C\in\Xi} C} \sum_{s\in S} u(f(s)).p(s)$. The result follows by Archimedean continuity and the affineness of u, as in the proof of part (i). □

Proof of Proposition 3. For part (i), by the argument in the proof of Theorem 3(i), for all $f \in \mathcal{A}$, $\underline{c_f} \preceq f$, and $u(\underline{c_f}) = \min_{p\in D(f,\underline{c_f})} \sum_{s\in S} u(f(s)).p(s)$. By B2, it follows that $\underline{c_f} \preceq^c f$. Moreover, for any $d \in \Delta(X)$ with $d \succ \underline{c_f}$, $d_{\frac{1}{2}}\underline{c_f} \succ \underline{c_f}$, whence, by the definition of $\underline{c_f}$, $d_{\frac{1}{2}}\underline{c_f} \not\preceq f$; so $d_{\frac{1}{2}}\underline{c_f} \succeq^c f$ by B4. Since $d \succ^c d_{\frac{1}{2}}\underline{c_f}$ by B1, it follows by A3 that $d \succ^c f$. By continuity, it follows that $\underline{c_f} \sim^c f$. Hence $\preceq^c$ is represented by $V(f) = u(\underline{c_f}) = \min_{p\in D(f,\underline{c_f})} \sum_{s\in S} u(f(s)).p(s)$ as required.

Similar arguments establish part (ii). □

B.3. Proofs of results in Section 5 and Appendix A

As stated in Section 5 (see in particular footnote 19), 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 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 B.12. *If the notion of stakes $\leqq^i_\bullet$ is monotone decreasing, 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 $((g_\alpha h)_\beta f, f) \leqq_{\preceq^i_{\Delta(X)}} \min\{(g, f), (h, f)\}$. By the representation (1) and the fact that the confidence ranking is balanced, it follows that $(g_\alpha h)_\beta f \succ^i f$. By 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 4. For any $x \in \Re^S_+$, let $\pi^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+,\ \text{if } z \succ x,\ \text{then } p.z > p.x\}$, and let $\bar{\pi}^i(x) = \{p \in \Delta(\Sigma) | \forall z \in \Re^S_+,\ \text{if } z \succ x,\ \text{then } p.z \geqslant p.x\}$. On inspection, it is straightforward to check that the reduced convexity property (Lemma B.12), combined with the concavity of u and the monotonicity of representation (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 (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 $p \in \bar{\pi}^i(x) \setminus \pi^i(x)$ for some i and x; so there exists z with $z \succ^i x$ and $p.z = p.x$. By the fact that stakes are monotone decreasing, the balancedness of the confidence ranking, and representation (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 $p.x = p.(\alpha z + (1-\alpha)x) > p.(z^i_\alpha x)$, 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, $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \setminus S_1)} > 0$; pick any $\delta > 0$ with $\min_{q \in \overline{\bigcup_{\mathcal{C} \in \Xi^i} \mathcal{C}}} \frac{q(S_1)}{q(S \setminus S_1)} > \delta \frac{\max_{s \notin S_1} u^{i'}(z_s)}{\min_{s \in S_1} u^{i'}(z_s)}$. 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 $p.(z + z^\epsilon) < p.x$ 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 $\bigcap_i \pi^i(x^i) \neq \varnothing$, 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 $\sum_s p(s)u^{i'}(x_s)(z_s - x_s) > 0$ for all $p \in \bigcap_{z \neq x} ri(D^i(x, z))$. Renormalising, it follows that, for any $q \in \Pi^i(f)$, $\sum_s q_s.(z_s - x_s) > 0$, and hence that $q.z > q.x$. 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 $\bar{p} \in \pi^i(x) \setminus \Pi^i(x)$. 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 $\bar{p}.y \leqslant b \leqslant q.y$ 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 $\alpha > 0$, $q.\alpha y = \frac{1}{\sum_{t \in S} p(t)u^{i'}(x_t)} \sum_s p(s)u^{i'}(x_s)\alpha y_s \geqslant 0$ 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 $\bar{p}.(x + \alpha y) \leqslant \bar{p}.x$, so $\bar{p} \notin \pi^i(x)$, as required. □

Proof of Proposition 4. We show this on Example 2. 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 (8) holds, namely:

$$\min\{\eta w \max\{|\delta - \bar{\delta}|, |\bar{\delta} - (1-\delta)|\}, 0.45\} + 0.5 > \frac{\bar{\delta}^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}}{\delta^{1-\gamma^1} - (1-\delta)^{1-\gamma^1}} \tag{B.2}$$

for all $\bar{\delta} < \nu$, then, for all $\bar{\delta}$ such that $x^2 \prec^2 z^2_{\bar{\delta}}$, $x^1 \not\prec^1 z^1_{\bar{\delta}}$; 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 5. 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:

$$\delta_2 < \delta_1^{0.5}\delta_2^{0.5} < \delta' < 1 - (1-\delta_1)^{0.5}(1-\delta_2)^{0.5} < \delta_1. \quad \text{(B.3)}$$

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 4 and the fact that the latter is a Pareto-improvement on the former, $\{(\frac{p(s_1)u^{1\prime}(x_j^1(s_1))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))}, \frac{p(s_2)u^{1\prime}(x_j^1(s_2))}{\sum_{t\in S} p(t)u^{1\prime}(x_j^1(t))}) \mid p \in ri(D^1(x_j^1, x_{j+1}^1))\} \cap \{(\frac{p(s_1)u^{2\prime}(x_j^2(s_1))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))}, \frac{p(s_2)u^{2\prime}(x_j^2(s_2))}{\sum_{t\in S} p(t)u^{2\prime}(x_j^2(t))} t) \mid p \in ri(D^2(x_j^2, x_{j+1}^2))\} = \varnothing$. 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:

$$\min\{\eta w \max\{|\delta_{j1} - \delta_{(j+1)1}|, |\delta_{(j+1)2} - \delta_{j2}|\}, 0.45\} + 0.5 \leqslant \frac{\delta_{j1} - \delta_{j1}\delta_{j2}}{\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2}}. \quad \text{(B.4)}$$

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, $0.95 > \frac{\delta_{j1}-\delta_{j1}\delta_{j2}}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}$. (B.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:

$$\delta_{(j+1)1} - \delta_{(j+2)2} \geqslant (\delta_{j1} - \delta_{j2})\left(1 - \frac{1}{\eta w}\frac{1}{\delta_{j1}+\delta_{j2}-2\delta_{j1}\delta_{j2}}\right). \quad \text{(B.5)}$$

But, using (B.3):

$$\delta_{j1} + \delta_{j2} - 2\delta_{j1}\delta_{j2} \geqslant \begin{cases} (1-\delta_1) + (1-\delta_1)^{0.5}(1-\delta_2)^{0.5}(2\delta_1 - 1) & \text{if } \delta_{j1}, \ \delta_{j2} > \frac{1}{2} \\ \delta_1^{0.5}\delta_2^{0.5} + \delta_2 - 2\delta_1^{0.5}\delta_2^{1.5} & \text{if } \delta_{j1}, \ \delta_{j2} < \frac{1}{2} \\ \frac{1}{2} & \text{if } \left(\delta_{j1} - \frac{1}{2}\right)\left(\delta_{j2} - \frac{1}{2}\right) \leqslant 0. \end{cases}$$

It thus follows from (B.5) that:

$$\delta_{(j+1)1} - \delta_{(j+1)2} \geqslant (\delta_{j1} - \delta_{j2})(1-\chi) \quad \text{(B.6)}$$

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 (B.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_{m+1}^1, x_{m+1}^2) is Pareto optimal. □

References

Abdellaoui, M., 2000. Parameter-free elicitation of utility and probability weighting functions. Manage. Sci. 46, 1497–1512.
Anscombe, F.J., Aumann, R.J., 1963. A definition of subjective probability. Ann. Math. Statist. 34, 199–205.
Aumann, R.J., 1962. Utility theory without the completeness axiom. Econometrica 30, 445–462.
Bewley, T.F., 1986/2002. Knightian decision theory. Part I. Decis. Econ. Finance 25, 79–110 (first version: 1986).
Bewley, T.F., 1989. Market innovation and entrepeneurship: A Knightian view, Tech. Rep. 905. Cowles Foundation.
Billot, A., Chateauneuf, A., Gilboa, I., Tallon, J.-M., 2000. Sharing beliefs: Between agreeing and disagreeing. Econometrica 68, 685–694.
Cerreia-Vioglio, S., Ghirardato, P., Maccheroni, F., Marinacci, M., Siniscalchi, M., 2011. Rational preferences under ambiguity. Econom. Theory 48, 341–375.
Cerreia-Vioglio, S., Giarlotta, A., Greco, S., Maccheroni, F., Marinacci, M., 2015. Rational Preference and Rationalizable Choice, Mimeo, Bocconi.
Chateauneuf, A., Faro, J.H., 2009. Ambiguity through confidence functions. J. Math. Econom. 45, 535–558.
Danan, E., 2003a. A behavioral model of individual welfare, Tech. rep. Université Paris 1.
Danan, E., 2003b. Revealed cognitive preference theory, Tech. rep. Université de Paris 1.
Danan, E., Gajdos, T., Hill, B., Tallon, J.-M., 2016. Robust social decisions. Amer. Econ. Rev. (forthcoming).
Dubra, J., Maccheroni, F., Ok, E.A., 2004. Expected utility theory without the completeness axiom. J. Econom. Theory 115, 118–133.
Faro, J.H., 2015. Variational Bewley preferences. J. Econom. Theory 157, 699–729.
Fishburn, P.C., 1970. Utility Theory for Decision Making. Wiley, New York.
Fon, V., Otani, Y., 1979. Classical welfare theorems with non-transitive and non-complete preferences. J. Econom. Theory 20, 409–418.
Galaabaatar, T., Karni, E., 2013. Subjective expected utility with incomplete preferences. Econometrica 81, 255–284.
Ghirardato, P., Maccheroni, F., Marinacci, M., 2004. Differentiating ambiguity and ambiguity attitude. J. Econom. Theory 118, 133–173.
Ghirardato, P., Siniscalchi, M., 2014. Risk-Sharing in the Small and in the Large. Mimeo Northwestern University.
Gilboa, I., Maccheroni, F., Marinacci, M., Schmeidler, D., 2010. Objective and subjective rationality in a multiple prior model. Econometrica 78.
Gilboa, I., Postlewaite, A., Schmeidler, D., 2009. Is it always rational to satisfy Savage's axioms? Econ. Philos. 285–296.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. J. Math. Econom. 18, 141–153.
Gul, F., Pesendorfer, W., 2006. Random expected utility. Econometrica 74, 121–146.
Hill, B., 2012. Confidence in preferences. Soc. Choice Welf. 39, 273–302.
Hill, B., 2013. Confidence and decision. Games Econom. Behav. 82, 675–692.
Hill, B., 2016. Confidence as a Source of Deferral, Mimeo HEC Paris.
Karni, E., Safra, Z., 2014. A Theory of Stochastic Choice under Uncertainty, Mimeo, Warwick Business School.
Klibanoff, P., Marinacci, M., Mukerji, S., 2005. A smooth model of decision making under ambiguity. Econometrica 73, 1849–1892.
Kopylov, I., 2009. Choice deferral and ambiguity aversion. Theor. Econom. 4, 199–225.
Lehrer, E., Teper, R., 2011. Justifiable preferences. J. Econom. Theory 146, 762–774.
Maccheroni, F., Marinacci, M., Rustichini, A., 2006. Ambiguity aversion, robustness, and the variational representation of preferences. Econometrica 74, 1447–1498.
Masatlioglu, Y., Ok, E.A., 2005. Rational choice with status quo bias. J. Econom. Theory 121, 1–29.
Minardi, S., Savochkin, A., 2015. Preferences with grades of indecisiveness. J. Econom. Theory 155, 300–331.
Nau, R.F., 1992. Indeterminate probabilities on finite sets. Ann. Statist. 20, 1737–1767.
Nau, R.F., 2006. The shape of incomplete preferences. Ann. Statist. 34, 2430–2448.
Nehring, K., 2009. Imprecise probabilistic beliefs as a context for decision-making under ambiguity. J. Econom. Theory 144, 1054–1091.
Ok, E.A., Ortoleva, P., Riella, G., 2012. Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes. Econometrica 80, 1791–1808.
Olofsgard, A., 2003. The political economy of reform: Institutional change as a tool for political credibility, Tech. rep.. World Bank.
Rigotti, L., Shannon, C., 2005. Uncertainty and risk in financial markets. Econometrica 73, 203–243.
Rigotti, L., Shannon, C., Strzalecki, T., 2008. Subjective beliefs and ex ante trade. Econometrica 76, 1167–1190.
Rodrik, D., 1996. Understanding economic policy reform. J. Econom. Lit. 9–41.
Schmeidler, D., 1971. A condition for the completeness of partial preference relations. Econometrica 39, 403–404.
Seidenfeld, T., Schervish, M., Kadane, J., 1995. A representation of partially ordered preferences. Ann. Statist. 23, 2168–2217.
Wakker, P.P., Deneffe, D., 1996. Eliciting von Neumann–Morgenstern utilities when probabilities are distorted or unknown. Manage. Sci. 42, 1131–1150.
Williamson, J., 1994. The Political Economy of Policy Reform. Peterson Institute.