# Confidence, aggregation and decision*

Brian Hill
CNRS & HEC Paris†

July 10, 2026

Dedicated to the memory of Philippe Mongin (1950-2020)

**Abstract**

Integrating experts' judgements into choice is a problem of belief aggregation for the purposes of subsequent decision. As a new approach to this problem, this paper develops an aggregation rule for confidence in beliefs, characterised by a novel Pareto condition that enjoins respecting consensuses borne of compromise. Confidence aggregation generalises standard probability aggregation rules such as linear pooling, while avoiding their spurious unanimity implications. It generates the first probability aggregation rule that can faithfully accommodate within-person cross-issue expertise diversity, hence resolving a longstanding challenge. Combined with existing accounts of decision under uncertainty, confidence aggregation yields a unified framework encompassing both Bayesian and non-Bayesian approaches to model misspecification, and provides a taxonomy that separates the belief aggregation and decision-making dimensions.

**Keywords:** Belief aggregation, model misspecification, confidence in beliefs, decision under uncertainty, linear pooling, spurious unanimity, expertise, consensus, model averaging.

**JEL codes:** D70, D81.

---

*Personal acknowledgements and thanks to be added.

†GREGHEC, 1 rue de la Libération, 78351 Jouy-en-Josas, France. E-mail: hill@hec.fr.

# 1 Introduction

How should policy decisions be taken on the the basis of honest and well-intentioned experts' beliefs? Conceptually, this involves belief aggregation for the purposes of subsequent decision. In practice, institutions at the science-policy frontier often reflect this two-stage structure. Science representatives such as the CDC Director in the US and the Chief Science Advisor in the UK are charged with providing policy-relevant assessments that reflect the range of expert opinion; institutions such as the IPCC aim to summarise scientific knowledge on a topic, in order to inform policy. They play the role of 'representatives', collating, evaluating and aggregating expert judgements, providing input into subsequent policy decisions. The same holds when the 'experts' are models. In fields ranging from economics to climate science, scientists draw on the output of informative yet imperfect models, together with their assessments of the uncertainty about those models' correctness, to form the judgements that they pass on to policy makers. As such, they act as representatives drawing on 'model expertise' to provide input for subsequent decision. These are instances of the *aggregation-for-subsequent-decision* problem, involving: a group of *individuals* each of whom provides input; a *representative* who aggregates them heeding her judgements about the individuals' respective expertise; and a *decision maker* who relies on the aggregate judgement to inform (policy) choice.

The classic Bayesian approach to this problem assumes the individuals' input to come in the form of probability measures; it aggregates them into a 'group' probability measure, which feeds into a Subjective Expected Utility (EU) decision rule. It applies whether the individuals are humans or models: indeed, the popular Bayesian Model Averaging approach to model uncertainty in statistics and economics (Raftery et al., 1997; Steel, 2020) essentially aggregates using a special case of linear pooling (e.g. Stone, 1961; Cooke, 1991; Gilboa et al., 2004)—i.e. the weighted average of probabilistic beliefs—with subsequent EU-based decision making.

However, this paradigm faces three apparently orthogonal challenges, concerning both the aggregation and the decision dimensions. This paper develops a single approach to the general aggregation-for-subsequent-

| | Labour | Real Estate | Both |
|---|---|---|---|
| Laura | 0.9 | 0.1 | 0.09 |
| Ray | 0.1 | 0.9 | 0.09 |
| Linear pool | $0.1 + 0.8w^L$ | $0.9 - 0.8w^L$ | 0.09 |

Table 1: Probability that a certain interest rate rise has a limited effect on the sector(s) in the top row

Final row gives the results of linear pooling, i.e. $p(E) = w^L p^L(E) + (1 - w^L)p^R(E)$ for every event $E$, with $w^L$ the weight for Laura, and $1 - w^L$ for Ray.

decision problem that resolves them all. We begin by presenting them.

The decision challenge lies in the purported inability of the Bayesian approach to properly accommodate fear of *model misspecification*. Approaches motivated by this concern typically take the same inputs—a probability measure for each individual—but inject them directly into a non-Expected Utility decision rule that allows for cautious attitudes towards model misspecification (Hansen and Sargent, 2001, 2022). However, since aggregation and decision are bundled into a single rule, it is difficult to disentangle whether the divergence from the Bayesian approach involves the decision stage, the aggregation stage, or both.

The two aggregation challenges can be illustrated on an example of a (two-member) central bank committee pondering whether to make a given interest rate rise. The committee agree that the determining factor in the choice is whether the rise has a limited (negative) effect on both the labour market and the real estate sector. Table 1 displays the two members' probability judgements for the rise having a limited effect on each of these sectors, and on both. Though each are competent economists, Laura is a specialist in the labour market, whilst Ray's field of expertise is the real estate sector. As is clear from the table, whilst they disagree significantly on the effect of the rise on each sector, they agree on the probability that it will have a limited effect on both sectors.

The first aggregation challenge concerns the treatment of unanimity by linear pooling, the result of which is shown in the final row of Table 1. Irrespective of the weights assigned to the individuals, the linear pool preserves their common judgement on the effect on both sectors—a consequence of the Pareto principle in this context (Mongin, 1995). However, the agreement on this probability is arguably *spurious*, resulting from the fortuitous interplay of two fundamental disagreements. After all, Laura gives a low probability to a limited effect on both sectors because of the low probability she assigns to the real estate event; Ray does so because of the low probability he assigns concerning the labour market; and they disagree on the judgements concerning labour and real estate alone. Several authors have argued that the automatic respect of such *spurious unanimity* is unjustified (Mongin, 2016), and hence a problem for linear pooling (Bradley, 2017b; Mongin and Pivato, 2020; Dietrich, 2021).

The second aggregation challenge involves the way that pooling rules typically incorporate expertise. They do so through the weights in the rule ($w^L$ in Table 1): each individual is allocated a single weight, with larger weights given to individuals with more expertise *overall*. They thus cannot reflect expertise differences across issues. For instance, linear pooling cannot capture the fact that Laura has more expertise on the labour market than the real estate sector (Genest and Zidek, 1986; French, 1985). However, in examples such as this, involving within-person expertise diversity, one might want to respect Laura's opinion more on labour and Ray's more on real estate. Linear pooling, like virtually all probability aggregation rules in the literature, does not allow this. Even model misspecification approaches (Hansen and Sargent, 2022; Cerreia–Vioglio et al., 2025) assign a single weight to each model, hence precluding expertise diversity within the same model across different issues (see Section 5.3).

All three challenges matter for decisions, as is clear in this example. Integration of model uncertainty or misspecification attitudes can clearly impact the decision. But even within EU, the aggregation issues affect the final decision: while linear pooling yields a low probability of a limited effect on both sectors, following each expert on their domain of expertise implies a substantially higher probability and thus potentially a different policy choice. Moreover, the decision-relevant event lies at the intersection of the committee members' fields of expertise, raising the question of how to combine their different levels of expertise across issues.

This paper develops a new approach to the general aggregation-for-subsequent-decision problem. By encompassing both Bayesian probability aggregation and recent model misspecification models, it situates the differences between them—hence tackling the decision challenge. It naturally integrates within-person cross-issue expertise diversity, whilst avoiding commitment to spurious unanimity—thus resolving the aggregation challenges.

The heart of the proposal, and first contribution of the paper, is a new aggregation rule for *confidence in beliefs*. Confidence in (experts' or models') beliefs is *prima facie* relevant for reflecting the extent of model uncertainty or misspecification, and several models of decision under uncertainty can be interpreted as involving confidence in beliefs in their representations (e.g. Marinacci, 2015; Maccheroni et al., 2006; Chateauneuf and Faro, 2009; Hill, 2013, 2019b; Bradley, 2017a). Our *confidence aggregation* rule enjoys independent normative justification: as we show, it is essentially characterised by a novel Pareto axiom. It thus preserves the Pareto intuition of respecting consensus, but appeals to a more refined notion of agreement that avoids the problem of spurious unanimity.

Our main contributions show how this single aggregation rule can overcome all the previous challenges in the context of existing approaches to decision under uncertainty. On the aggregation front, popular probability aggregation rules can be recovered as special cases of confidence aggregation, corresponding to particular assumptions on confidence in the individuals' beliefs. By identifying the point of divergence with classic pooling rules, this analysis sets the stage for the integration of within-person cross-issue expertise diversity. If an individual has more expertise on one issue than another, that justifies more confidence *ceteris paribus* in her beliefs concerning the former issue. Drawing on this insight, we explore the consequences of confidence aggregation when applied in cases involving different degrees of confidence—reflecting differing expertise—across different issues. It yields aggregate judgements that more strongly respect an individual's judgement on the issues on which she has more expertise, as compared to those on which she has less. The approach thus resolves the expertise challenge; moreover, since it does not respect spurious unanimities resulting from ignoring expertise differences, it tackles that aggregation challenge too.

As an application, we use confidence aggregation to generate a new family of probability aggregation rules that can accommodate within-person expertise diversity. To our knowledge, these *expertise-sensitive pooling rules* are the first of their sort in the literature.

On the decision front, known preference axioms connect the confidence aggregation rule to existing decision approaches. These include the confidence-in-beliefs decision framework (Hill, 2013, 2019b), a generalisation of variational preferences (Maccheroni et al., 2006), which include in turn the main model misspecification rules. This provides a general model of aggregation-for-subsequent-decision on the basis of multiple individual inputs. It encompasses all the aforementioned approaches—from linear pooling and Bayesian Model Averaging to model misspecification—situating them along three dimensions. Two of these dimensions—reflecting assumptions about confidence in individual judgements and the only parameter in our aggregation rule—are concerned with the aggregation step; the last—reflecting the importance of the decision, and hence confidence required to take it—concerns decision. As such, our rule situates the divergences and communalities among various existing approaches, as well as suggesting new ones. For instance, combining the techniques used to develop the new expertise-sensitive pooling rule with the settings yielding model misspecification representations, we generate novel expertise-sensitive model misspecification decision rules.

The paper is organised as follows. Sections 2 and 3 focus entirely on aggregation: the former sets out the framework and the aggregation rule, whereas the latter shows how it overcomes the aggregation challenges. Preferences are introduced in Section 4, which contains a preference-based axiomatisation of confidence aggregation. Section 5 expands it to a decision framework encompassing Bayesian as well as model misspecification approaches. Section 6 discusses remaining related literature. Proofs and supplementary material are contained in the Appendices.

# 2 Confidence aggregation

The general aggregation-for-subsequent-decision problem involves: a set $I$ of *individuals*, each of which provides input in a specified format; a *representative* who aggregates this input taking into account her evaluations of the individuals' expertise; a *decision maker* who relies on the representative's aggregate judgements to take the subsequent decision(s).

The representative and decision maker will ultimately be modelled by preferences (Sections 4 and 5). By contrast, we make no assumption on the format of individuals' inputs. For instance, they could come in the form of probability measures, as in probability aggregation or model misspecification approaches (Genest and Zidek, 1986; Cerreia–Vioglio et al., 2025). Alternatively, they could be preferences, as in the literature on preference aggregation. As such, the development here can support several interpretations, some of which will be highlighted presently.

This section and the subsequent one will focus on the aggregation stage, involving the individuals and the representative. This section sets out the heart of our approach: a confidence aggregation rule that takes a representation of confidence in belief for each individual as input, and provides a single confidence-in-belief representation as output. *Inter alia*, Section 3 illustrates how it can apply to the general problem (involving alternative input formats).

## 2.1 Preliminaries

**Setup** Let $\Omega$ be a non-empty set of *states* with a $\sigma$-algebra $\Sigma$ of subsets, called *events*. *Partitions* are sets of mutually disjoint events whose union is $\Omega$. For any partition $\mathcal{P}$ (including $\Omega$ itself), $\Delta(\mathcal{P})$ denotes the set of probability measures over $\mathcal{P}$; henceforth, we let $\Delta = \Delta(\Omega)$, and assume it to be compact.[1] For any $p \in \Delta$ and partition $\mathcal{P}_j$, $p|_{\mathcal{P}_j} \in \Delta(\mathcal{P}_j)$ denotes the projection of $p$ onto $\Delta(\mathcal{P}_j)$.

$O \subseteq \mathbb{R}_{\geqslant 0}$ with $0 \in O$ is an ordered set of confidence levels, endowed with the (strict) order $>$ inherited from $\mathbb{R}$. $\geqslant$ is the corresponding weak order. No general assumptions will be made about the cardinality of $O$ in this paper: we only assume that, if $O$ is not finite, then it is an interval in

[1] Throughout, we take the weak* topology on $\Delta$ and $\Delta(\mathcal{P})$. Most of the development in this paper holds for both finitely and countably additive probabilities, though some examples make more sense with one rather than the other. So, for instance, $\Delta$ is automatically compact if $\Omega$ is finite or $\Delta$ is the set of finitely additive probabilities (by the Banach-Alaoglu Theorem); when $\Delta$ is the space of countably additive probabilities, compactness requires further assumptions, e.g. that $\Omega$ is compact Polish space. When the difference matters for results and is not clear from the context, the character of the measures will be specified.

$\mathbb{R}$, with the associated topology.[2] We use vector notation to denote tuples of confidence levels, i.e. elements of $O^n$ such as $\mathbf{o} = (o_1, \ldots, o_n)$. With slight abuse of notation, we use $\geqslant$ to denote the dominance relation on such tuples: $\mathbf{o} \geqslant \mathbf{o}'$ if and only if $o_i \geqslant o'_i$ for all $i = 1, \ldots, n$.

**Beliefs and confidence** We work with a general model of confidence in beliefs that, as explained below, underlies many recent models of decision under uncertainty (Hill, 2019b). The belief state of an agent, incorporating confidence, is represented by a *confidence ranking*, defined as follows.

**Definition 1.** A *confidence ranking* is an upper semicontinuous, increasing function $c : O \cup \{\infty\} \to 2^{\Delta} \backslash \varnothing$ with $c(\infty) = \Delta$.[3]

The *centre* of confidence ranking $c$ is its smallest element, $c(0)$. Confidence ranking $c$ is *pointed* if its centre is a singleton. It is *convex* (respectively, *closed*) if, for every $o \in O$, $c(o)$ is a convex (resp. closed) set.

The *implausibility function* $\iota : \Delta \to O \cup \{\infty\}$ generated by confidence ranking $c$ is defined by $\iota(p) = \min \{o \in O : p \in c(o)\}$ whenever the set is non-empty, and $\iota(p) = \infty$ otherwise.

For each confidence level $o$, the set of priors $c(o)$ picked out by the confidence ranking $c$ represents the beliefs the agent holds with confidence at least $o$. Identifying a probability judgement—such as 'the probability of event $E$ is greater than $x$'—with the set of probability measures where it holds, a judgement $\mathcal{J} \subseteq \Delta$ is held with confidence at least $o$ if it contains the corresponding set in the agent's confidence ranking, i.e. $\mathcal{J} \supseteq c(o)$. Confidence rankings may or may not have singleton sets at their centre. Those that do represent Bayesians with confidence: agents who assign a precise probability to every event (that given by the centre), though may have more confidence in some judgements than others (as represented by the rest of the confidence ranking).

For any $O' \subseteq O$ with $o' \in O'$ whenever $o' \geqslant o''$ and $o'' \in O'$, and any increasing, upper semicontinuous function $c : O' \to 2^{\Delta} \backslash \varnothing$, the *natural extension* of $c$, denoted $\bar{c}$, is the confidence ranking defined by $\bar{c}(o') = c(o')$

[2]It follows that $\geqslant$ is continuous: its upper and lower contour sets are closed.

[3]I.e. for all $o \geqslant o'$, $c(o) \supseteq c(o')$ and for any decreasing sequence $o_i \in O$ with $o_i \to o$, $c(o) = \bigcap_i c(o_i)$.

for $o' \in O'$ and $\overline{c}(o') = \bigcap_{o'' \in O'} c(o'')$ otherwise. The *convex closure* $c^{clconv}$ of a confidence ranking $c$ is defined in the natural way: for all $o \in O$, $c^{clconv}(o) = clconv(c(o))$, where $clconv(X)$ for a set $X \subseteq \Delta$ is the closure of the convex hull of $X$. $\Pi$ denotes the set of confidence rankings.

The final clause of Definition 1 provides an alternative, equivalent representation of confidence rankings as *implausibility functions*. They yield the 'implausibility' of each probability measure, in terms of the smallest confidence level at which the probability measure doesn't contradict any judgement held with that much confidence.[4]

**Related models and distance-based confidence** The representation of confidence in beliefs used here includes that in Hill (2013, 2019b) under the calibration developed in Hill (2019a) (see Section 6). Given its equivalence to real-valued implausibility functions on the probability space, it also underlies models of decision under uncertainty which involve such functions in their preference representation. Prominent examples include smooth, variational, multiplier and confidence preferences (Klibanoff et al., 2005; Maccheroni et al., 2006; Hansen and Sargent, 2001; Chateauneuf and Faro, 2009), some of which interpret such functions in terms of confidence.

Moreover, this alternative implausibility-function representation implies that confidence rankings can be generated from single probability measures and (statistical) distances. A *statistical distance* $\rho$ on $\Delta(\mathcal{P})$ is a function $\rho : \Delta(\mathcal{P})^2 \to [0, \infty]$ such that: $\rho(p, q) = 0$ if and only if $p = q$; and $\rho(\bullet, q)$ is a lower semicontinuous function, for every $q \in \Delta(\mathcal{P})$. A distance $\rho$ is *convex* if, for every $q \in \Delta(\mathcal{P})$, the function $\rho(\bullet, q)$ is strictly convex on $\{p \in \Delta(\mathcal{P}) : \rho(p, q) \in [0, \infty)\}$.[5] A (convex) *classical statistical distance* $d$ is the specification, for each partition $\mathcal{P}$ (including $\Omega$ itself), of a (convex) statistical distance on $\Delta(\mathcal{P})$. With slight abuse of notation, we use $d$ to refer to the distance for each $\Delta(\mathcal{P})$.

The confidence ranking with centre at probability $p \in \Delta$ and generated by distance $\rho$ on $\Delta$ and weight $w \in \mathbb{R}_{>0}$ is defined as follows.

[4]Note that $c$ can be defined from $\iota$: $c(o) = \{p \in \Delta : \iota(p) \leqslant o\}$. It follows immediately that the implausibility function $\iota$ is lower semicontinuous if $c$ is closed.

[5]That is, for all $p, r \in \Delta(\mathcal{P})$ with $\rho(p, q), \rho(r, q) \neq \infty$, $p \neq r$ and $\alpha \in (0, 1)$, $\rho(\alpha p + (1 - \alpha)r, q) < \alpha\rho(p, q) + (1 - \alpha)\rho(r, q)$.

| Generating distance | $\rho(q,p) =$ | $w$ $\rho$-confidence ranking generated by $p$, $c(o) =$ |
|---|---|---|
| Euclidean | $\sum_{\omega\in\Omega}(q(\omega)-p(\omega))^2$ | $\{q \in \Delta : w\sum_{\omega\in\Omega}(q(\omega)-p(\omega))^2 \leqslant o\}$ |
| Relative entropy | $R(q\|p)$ | $\{q \in \Delta : wR(q\|p) \leqslant o\}$ |
| Reverse relative entropy | $R(p\|q)$ | $\{q \in \Delta : wR(p\|q) \leqslant o\}$ |
| Gini | $G(q\|p)$ | $\{q \in \Delta : wG(q\|p) \leqslant o\}$ |

Table 2: Examples of classical statistical distances and distance-generated confidence rankings. The distances in the first, second and fourth rows are convex.

Note: Euclidean distance only well-defined on finite $\Omega$. $R$ is the relative entropy, defined by: $R(p\|q) = \int_\Omega \left(\log \frac{dp}{dq}(\omega)\right) dp(\omega)$ for $p$ absolutely continuous with respect to $q$, and $\infty$ otherwise. $G$ is the Gini divergence, defined by $G(p\|q) = \int_\Omega (\frac{dp}{dq}(\omega) - 1)^2 dq(\omega)$ for $p$ absolutely continuous with respect to $q$, and $\infty$ otherwise.

**Definition 2.** Let $p \in \Delta$ be a probability measure, $\rho$ a statistical distance on $\Delta$ and $w \in \mathbb{R}_{\geqslant 0}$. The *$w$ confidence ranking generated by $p$ under $\rho$*—or simply the *$w$ $\rho$-confidence ranking generated by $p$*—is defined by $c(o) = \{q \in \Delta : w\rho(q,p) \leqslant o\}$ for all $o \in O$.

This definition is most natural under $O = \mathbb{R}_{\geqslant 0}$, which will henceforth be assumed in discussions featuring these confidence rankings. The $w$ $\rho$-confidence ranking generated by $p$ is clearly associated with implausibility function $\iota(q) = w\rho(q,p)$. Table 2 lists some well-known classical statistical distances, and the corresponding generated confidence rankings. Several popular models of decision under uncertainty use such distance-generated confidence rankings; for instance, the $w$ relative-entropy-confidence ranking is involved in multiplier preferences (Hansen and Sargent, 2001).

## 2.2 Consensus-preserving confidence aggregation

A confidence aggregation rule transforms a profile of confidence rankings into a confidence ranking, where a *profile* is a tuple $(c^1, \ldots, c^n)$ of confidence rankings. In this section, we refer to the confidence rankings in the profiles as those of the individuals, and the resulting one as that of the group's representative; alternative interpretations will be discussed in Section 3.

The central motivation for the aggregation rule developed here is, like the Pareto principle, to respect consensus; unlike standard Pareto, however, it recognises that consensus is the product of compromise. To this end, note that each tuple $\mathbf{o} = (o_1, \ldots, o_n)$ of confidence levels generates a tuple of sets of priors $(c^1(o_1), \ldots, c^n(o_n))$, reflecting the beliefs held by each individual at the confidence level specified for him. If the intersection of these sets is empty, $\bigcap_i c^i(o_i) = \varnothing$, these beliefs are in contradiction. If $\bigcap_i c^i(o_i) \neq \varnothing$, they are not: there is a consistent overall *consensus* position, characterised by $\bigcap_i c^i(o_i)$, which incorporates the beliefs of each individual at the assigned confidence level. In this consensus, each individual $i$ only puts beliefs held with confidence $o_i$ or more 'on the table', 'setting aside' lower-confidence beliefs. All the former beliefs are retained in the consensus,[6] while the individuals compromise by 'bracketing out' the latter ones. To that extent, the consensus is borne of compromises, which are regulated by the confidence levels determining the beliefs each individual contributes.

There may thus be several such consensuses, involving different compromises—i.e. different tuples of confidence levels. To translate them into levels of confidence deemed relevant for the group, we use a *confidence-level aggregator*: an operator $\otimes : (O \cup \{\infty\})^n \to O \cup \{\infty\}$ that is monotonic in each argument, i.e. such that for every pair of profiles of confidence levels $\mathbf{o}, \mathbf{o}'$ with $\mathbf{o} \geqslant \mathbf{o}'$, $\otimes\mathbf{o} \geqslant \otimes\mathbf{o}'$. For a consensus corresponding to individual confidence levels $\mathbf{o}$, the confidence-level aggregator picks out the group confidence warranted in the associated consensus judgements. Monotonicity means that the higher the individual confidence levels $\mathbf{o}$ behind the consensus, the higher the corresponding confidence level for the group's representative. This reflects the fact that higher individual confidence levels translate into a more robust consensus, albeit one involving more compromise.

In our preference-based characterisation (Section 4), the relevant confidence-level aggregator will be endogenous; for illustration, Table 3 provides some examples.

[6] Formally, for a probability judgement $\mathcal{J}$, if some individual $i$ holds it at this level of confidence (i.e. $c^i(o_i) \subseteq \mathcal{J}$), then it holds in the consensus ($\bigcap_i c^i(o_i) \subseteq \mathcal{J}$).

| **Aggregator** | **Definition** |
|---|---|
| Average aggregator | $\otimes\mathbf{o} = \sum_{i=1}^{n} \frac{1}{n} o_i$ |
| Maximum aggregator | $\otimes\mathbf{o} = \max\{o_i\}$ |
| Minimum aggregator | $\otimes\mathbf{o} = \min\{o_i\}$ |

Table 3: Examples of confidence-level aggregators.

Each confidence-level aggregator generates a confidence aggregation rule—a function from profiles of confidence rankings to confidence rankings—as defined below. Since in our preference-based characterisation (Section 4) we follow the economic literature and work in a single-profile setup, we also provide the associated definition of consensus-preserving confidence aggregation for a fixed confidence ranking $c^{rep}$ and profile $(c^1, \ldots, c^n)$.

**Definition 3.** Let $\otimes : (O \cup \{\infty\})^n \to O \cup \{\infty\}$ be a confidence-level aggregator. The *consensus-preserving confidence aggregation rule* $F_\otimes : \Pi^n \to \Pi$ is such that, for every profile $(c'^1, \ldots, c'^n) \in \Pi^n$, $F_\otimes(c'^1, \ldots, c'^n) = \overline{\Phi_\otimes(c'^1, \ldots, c'^n)}$, where, for every $o \in O$ such that $\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o} \bigcap_i c'^i(o_i) \neq \varnothing$

$$\Phi_\otimes(c'^1, \ldots, c'^n)(o) = \bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o} \bigcap_{i=1}^{n} c'^i(o_i). \tag{1}$$

For a fixed confidence ranking $c^{rep} \in \Pi$ and profile of confidence rankings $(c^1, \ldots, c^n) \in \Pi^n$, $c^{rep}$ is a *consensus-preserving confidence aggregation of* $(c^1, \ldots, c^n)$ if there exists a confidence-level aggregator $\otimes$ such that $c^{rep} = F_\otimes(c^1, \ldots, c^n)$. In this case, we say that $c^{rep}$ is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$ *under* $\otimes$.

Under consensus-preserving confidence aggregation—or *confidence aggregation* for short—the aggregate judgements held with confidence level $o$ are formed by looking at the consensuses considered to warrant confidence level $o$ or less according to $\otimes$.[7] More specifically, the group's represen-

[7] The use of consensuses corresponding to confidence levels less than or equal to $o$ in Eq. (1) ensures that the image of $F_\otimes$ is a well-defined confidence ranking, no matter the $\otimes$. As discussed in Appendix B, 'less than or equal to' can be replaced by 'equal to' for

tative holds a probability judgement with confidence $o$ if that judgement holds for all such consensuses: this is guaranteed by the union in Eq. (1). Judgements that hold unanimously across the appropriate consensuses are thus preserved. However, if two consensuses contradict each other on an issue, neither's judgement will be retained by the representative at that confidence level.

Note that this aggregation rule implies that the representative's and individuals' confidence in a judgement co-vary: the former's confidence in a judgement is higher when the individual beliefs feeding into the relevant consensuses are held with more confidence.

Illustrative applications of confidence aggregation will be provided in the next Section. For reference, note that, as shown in Proposition D.1 (Appendix D.1), this aggregation rule can be formulated in terms of implausibility functions: $c^{rep}$ is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$ under $\otimes$ if and only if, for all $p \in \Delta$

$$\iota^{rep}(p) = \begin{cases} 0 & p \in \arg\min_{p' \in \Delta} \otimes(\iota^1(p'), \ldots, \iota^n(p')) \\ \otimes(\iota^1(p), \ldots, \iota^n(p)) & \text{otherwise.} \end{cases} \tag{2}$$

# 3 Confidence, probability aggregation and expertise diversity

Input into aggregation-for-subsequent-decision problems often does not come in the form of confidence rankings: a more common format is as probability measures. In this section, we consider how confidence aggregation can deal with the challenges for existing approaches to aggregating probabilities. We first show that standard pooling rules can be recovered as special cases of confidence aggregation. Analysing the underlying assumptions will naturally reveal how it can resolve both of the aggregation challenges mentioned in the Introduction. The discussion also contains several examples illustrating confidence aggregation, and culminates in a new probability aggregation rule integrating within-person cross-issue expertise diversity.

all $\otimes$ in Table 3.

$$
\begin{tikzcd}
\text{Tuple of conf. rankings } \mathbf{c} \arrow[r, "\textit{conf. aggregation}"] & \text{Conf. ranking } c \arrow[d, "\textit{centre}"] \\
\text{Tuple of probabilities } \mathbf{p} \arrow[u, "\textit{generate}"] \arrow[r, Rightarrow, "\textit{derived pbty aggregation}"] & \text{Probability } p
\end{tikzcd}
$$

Figure 1: Using confidence aggregation to generate probability aggregation rules

### 3.1 Recovering probability aggregation from confidence aggregation

Probability aggregation takes as input a profile of probability measures $\mathbf{p} = (p_1, \ldots, p_n) \in \Delta^n$. To connect pooling rules operating on such profiles with confidence aggregation, recall that, once a (statistical) distance and weight are specified, each probability measure generates a unique pointed confidence ranking (Definition 2, Section 2.1). This provides the following possibility for using consensus-preserving confidence aggregation (Section 2.2) to aggregate probability measures. Given a profile of probability measures, take the profile of confidence rankings generated by them under a given distance. Picking a confidence-level aggregator, confidence aggregation can be applied on them, to produce a confidence ranking, call it $c$. This naturally identifies the 'best-guess' set of probability measures, namely $\min_{o \in O} c(o)$. If this is a singleton (i.e. $c$ is pointed), then the procedure yields a unique probability measure, as required by probability aggregation rules. This schema is summarised in Figure 1.

The following result compares this probability aggregation method to standard pooling rules.

**Proposition 1.** *Let $\mathbf{p} = (p_1, \ldots, p_n) \in \Delta^n$ be a profile of probability measures, and $(w^1, \ldots, w^n)$ an tuple of weights, with $w^i \geqslant 0$ for all $i$, with strict inequality for some $i$. For each row in Table 4 for which the pooling rule in the second column is defined, the following holds:*

*(*) Let $c$ be the consensus-preserving confidence aggregation under the average confidence-level aggregator of $w^i$ $\rho$-confidence rankings gen-*

| Generating distance $\rho$ | Pooling rule | Centre $p$ satisfies |
|---|---|---|
| Euclidean | Linear pooling | $p = \sum_i \frac{w^i}{\sum_{i=1}^n w^i} p_i$ |
| Relative entropy | Geometric pooling | $p(\omega) \propto \prod_i p_i^{\frac{w^i}{\sum_{i=1}^n w^i}}(\omega)$ |
| Reverse relative entropy | Linear pooling | $p = \sum_i \frac{w^i}{\sum_{i=1}^n w^i} p_i$ |
| Gini | Harmonic pooling | $p(\omega) \propto \left(\sum_i \frac{w^i}{p_i(\omega)}\right)^{-1}$ |

Table 4: The pooling rules derived from confidence aggregation applied to confidence rankings generated under given classical distances (as in Figure 1). To be read in the context of Proposition 1.

*erated by $p_i$, where $\rho$ is the distance given in the first column of Table 4. Then its centre is a singleton containing the pool of the $p_i$ under the rule specified in the second column of the Table, with weights $\frac{w^i}{\sum_{i=1}^n w^i}$. In other words, the centre contains the probability measure satisfying the equation in the third column of the Table.*

The first three rows of Table 4 show that the two most prominent pooling rules in the literature (Genest and Zidek, 1986; Mongin, 1995; Dietrich, 2021) correspond to special cases of confidence aggregation. Figure 2b provides a graphical illustration of this result on the example from the Introduction, which will be further analysed below (Example 3.1). Central to it is the use of specific classical statistical distances (Table 2) to generate confidence rankings from individuals' probability measures. As is clear from the comparison of the cases in Table 4, the 'shape' of the confidence rankings determines the pooling rule reproduced. In this sense, the use of, say, linear pooling involves an assumption leading to the use of confidence rankings generated by the Euclidean or reverse relative entropy distances. And the evaluation of this pooling rule can thus pass via an appraisal of such assumptions.[8]

The $w^i$ $\rho$-confidence rankings in Proposition 1—and hence the as-

[8]Given that, as noted in Section 2.1, a distance and a probability measure generate a confidence ranking, Proposition 1 is technically related to a literature characterising aggregation rules in terms of distances in probability space (e.g. Abbas, 2009; Kemeny, 1959 initiated a similar approach for preference aggregation). This literature takes the distances as given, whereas we consider them as purported representations of belief states—and, as shall be clear below, evaluate them as such.

sumptions underlying them—support several interpretations. Under the *representative's-confidence* interpretation, they reflect the opinions the representative forms on the basis of each individual's input taken separately. For instance, if the input probability distributions come from models, each confidence ranking reflects the representative's assessment of the quality of that model; and similarly for Bayesian experts. Indeed, as is clear in Proposition 1, the weights in the generated pooling rules, which are often interpreted in terms of the representative's evaluation of individuals' expertise (e.g. McConway, 1981), are generated by the confidence rankings. Interpreted this way, the use of confidence rankings generated by a specific distance assumes that they faithfully capture the representative's confidence assessments given each individual's input. If they don't reflect the beliefs she infers from the provided probabilities and her evaluation of the individuals' expertise, then the postulated distance is inappropriate.

Under an alternative *individuals'-confidence* interpretation, the individuals can assess their confidence in their probability judgements; for instance, they may be experts whose belief states are richer than the probability measures requested. Under this interpretation, consensus-preserving confidence aggregation could be applied on the individuals' full beliefs (including confidence); it is the setup of the probability aggregation problem that artificially restricts the input to probability measures. Compared to this 'full-belief' benchmark, using confidence rankings generated by a specific distance assumes that they properly represent individuals' actual confidence in their beliefs. The assumption will be unjustified if individuals' confidence is not well reflected by the distance used.

All of the results and points made in this paper hold under both of these interpretations (and indeed, under hybrid ones where the stipulated confidence reflects individuals' confidence and the representative's assessment). Nevertheless, to lighten the discussion, we shall often conduct it in terms of the latter interpretation.

Proposition 1 suggests a strategy for tackling the aggregation challenges cited in the Introduction. If traditional pooling rules follow from specific assumptions on confidence, then scrutinizing these assumptions may provide the key to overcoming their weaknesses.

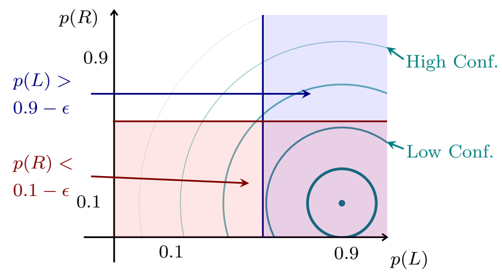


(a) Illustration of 'issue-neutrality'.
**Note:** The blue area represents the probability judgement, $\mathcal{L}_\epsilon$, that $p(L)$ is within $\epsilon$ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement, $\mathcal{R}_\epsilon$, that $p(R)$ is within $\epsilon$ of $p^L(R) = 0.1$. The confidence in these judgements (corresponding to the largest circular set contained in each area; Section 2.1) is the same.

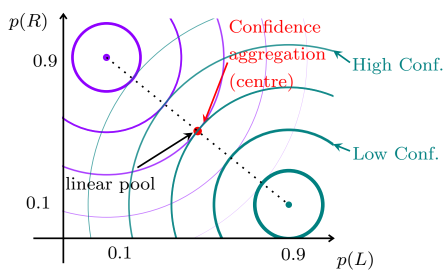


(b) Illustration of Proposition 1.
**Note:** The red point is the centre of the result of confidence aggregation applied to the two confidence rankings (Proposition 1). Each point on the dotted line is obtained by linear pooling (with some choice of weights). This graph displays the case of $w^L = w^R$; other cases produce centres lying on the dotted line (i.e. coinciding with some linear pool).

Figure 2: Confidence rankings generated as in Proposition 1.
**Note:** Each graph shows the space of pairs of probability values $(p(L), p(R))$ for the Labour and Real Estate events ($L$ and $R$; Example 3.1). The areas (sets of probability values) enclosed by the green circles represent the $w^L$ Euclidean confidence ranking generated by Laura's probability $p^L$ (Definition 2): they are the projection of the confidence ranking into this space. Larger, lighter circles correspond to higher confidence levels. The purple circles represent the $w^R$ Euclidean confidence ranking generated by $p^R$ (Ray's probabilities), with $w^R = w^L$.

## 3.2 Classical statistical distances and expertise

One specificity of the classical-distance-based confidence rankings involved in Proposition 1 is a certain 'neutrality' to the identity of the issues involved. All that counts for the confidence with which a probability judgement is held is the classical distance from the centre to the closest probability measure where the judgement doesn't hold—independently of the issue concerned by the judgement. We illustrate this on a running example.

**Example 3.1.** Formalise the example from the Introduction with a state space $\Omega = \{\omega_{LR}, \omega_L, \omega_R, \omega_N\}$ where $\omega_{LR}$ (respectively $\omega_L$, $\omega_R$, $\omega_N$) is the state in which there is a limited effect on both the labour and real estate sectors (resp. only the labour market, only the real estate sector, neither). So the event that there is a limited effect on the labour market is $L = \{\omega_{LR}, \omega_L\}$; the corresponding event for real estate is $R = \{\omega_{LR}, \omega_R\}$. Laura's probability judgements (Table 1) define the measure $p^L$ with $p^L(\omega_{LR}) = 0.09$, $p^L(\omega_L) = 0.81$, $p^L(\omega_R) = 0.01, p^L(\omega_N) = 0.09$.[9] Hence, for any $\epsilon \in [0, 0.9]$, she holds both the judgement, $\mathcal{L}_\epsilon$, that the probability of $L$ is greater than $0.9 - \epsilon$, and the judgement, $\mathcal{R}_\epsilon$, that the probability of $R$ is less than $0.1 + \epsilon$.[10] Note that these judgements involve moving the same amount away from her best-guess probability for $L$ (0.9) and $R$ (0.1) respectively. Which of them is she more confident in?

Proposition D.3 (Appendix D.2) shows that, if her confidence ranking is generated by any of the classical statistical distances in Table 4, the confidence in the two judgements is the same, for all $\epsilon$. Figure 2a illustrates the intuition: given the 'rotational symmetry' of the sets of priors in the confidence ranking, the highest confidence levels at which the judgements hold are the same. Hence the confidence assigned to a judgement that 'deviates' from the best-guess probability by a certain amount depends, in this example, only on the extent of the deviation, not on the issue concerned by the judgement.

The confidence rankings generating standard pooling rules thus represent individuals as having or warranting the same confidence in the probability judgements encoded in their probability measure, regardless of the issues that these judgements concern. As such, they cannot properly capture an individual with different confidence in judgements pertaining to different issues. The previous example is arguably such a case. Recall that Laura has more expertise on one issue (labour) than another (real estate). But an expertise difference typically translates into a difference in confidence: *ceteris paribus* she will have more confidence in her judgements concerning her issue of expertise than in those that do not. In other words, the confidence rankings involved in Proposition 1, based on classical statistical distances on the probability space, assume that there is no

[9] Throughout the discussion of the example in this Section, superscripts will indicate the individual ($L$ for Laura; $R$ for Ray). To aid legibility, there are coloured using the same colour scheme as in Figures 2 and 3.

[10] I.e. $\mathcal{L}_\epsilon = \{p \in \Delta : p(L) \geqslant 0.9 - \epsilon\}$ and $\mathcal{R}_\epsilon = \{p \in \Delta : p(R) \leqslant 0.1 + \epsilon\}$.

within-person cross-issue difference in expertise.

This observation brings a new perspective on the problem that linear pooling and other standard pooling rules have with within-person expertise diversity. The source of the problem isn't so much the underlying rule in our reconstruction—confidence aggregation—but the use of confidence-ranking-generating procedures which *de facto* assume away within-person cross-issue expertise differences. It thus suggests that confidence aggregation applied to confidence rankings that *do* correctly capture expertise differences could incorporate more faithfully such differences into the resulting group beliefs. To confirm this suggestion, we first show how to represent expertise differences with confidence rankings.

## 3.3 Representing expertise using confidence rankings

For the presentation, let us focus on issues that can be related to events in $\Omega$; see Appendix A.2 for a generalisation. Consider a sequence $\mathcal{P}_1, \dots \mathcal{P}_m$ of partitions of $\Omega$; each partition could be thought of as an *issue*. For instance, a partition could just be an event $E$ and its complement: the issue is whether the event holds. Another partition could have cells corresponding to the event that a parameter takes a given value: the issue is the value of the parameter. We say that a sequence of partitions $\mathcal{P}_1, \dots \mathcal{P}_m$ is *rich* if, for any $(p_1, \dots, p_m) \in \prod_{j=1}^m \Delta(\mathcal{P}_j)$, there exists at most one $p \in \Delta$ with $p|_{\mathcal{P}_j} = p_j$ for all $j = 1, \dots, m$. When the sequence of partitions is rich, then each tuple of probability measures, one on each partition, determines at most one probability measure over the whole space.

**Example 3.2.** In the example from the Introduction, with the state space and events defined in Example 3.1, each of the three issues in Table 1 corresponds to a two-element partition: $\mathcal{P}_L = \{L, L^c\}$ (whether there will be an effect on the labour market), $\mathcal{P}_R = \{R, R^c\}$ (concerning real estate), $\mathcal{P}_B = \{B, B^c\}$, where $B = \{\omega_{LR}\} = L \cap R$ (whether there will be an effect on both). Clearly every specification of a probability on each of these partitions determines at most one probability on $\Omega$, so this sequence of partitions is rich.

Now consider the following family of confidence rankings.

**Definition 4.** Let $\mathcal{P}_1, \ldots \mathcal{P}_m$ be partitions of $\Omega$ and $d$ be a classical statistical distance. For any probability measure $p \in \Delta$, and any vector $\mathbf{w} = (w_1, \ldots, w_m)$ of positive real-valued weights, the $\mathbf{w}$ *$d$-confidence ranking generated by $p$* is defined as: for each $o \in O$,

$$c(o) = \left\{ q \in \Delta : \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \leqslant o \right\} \tag{3}$$

For such confidence rankings, at each confidence level, the corresponding set of priors are those for which the weighted sum of the distances from the centre probability, taken over all the partitions (or issues), is less than a certain value. These belong to the family of distance-generated confidence rankings in Definition 2, with the following implausibility function:

$$\iota(q) = \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \tag{4}$$

Apart from the special case involving a single partition $\mathcal{P} = \Omega$, such $\iota$ are not (multiples of) classical statistical distances (Section 2.1).[11]

The issue-specific weights in $\mathbf{w}$ $d$-confidence rankings can capture relative expertise across issues, with higher weights on a given issue translating more confidence in judgements concerning it. This can be seen on a continuation of the running example.

**Example 3.3.** Consider $p^L$ as in Example 3.1, and consider the confidence ranking generated by it with Euclidean distance and vector of weights $\mathbf{w^L} = (w_L^L, w_R^L, w_B^L)$, namely:

$$c^L(o) = \left\{ q \in \Delta : \sum_{j=\{L,R,B\}} w_j^L (q(j) - p^L(j))^2 \leqslant o \right\} \tag{5}$$

The weights reflect the relative confidence in judgements about $L$, $R$ and $B$. Larger weights involve a higher 'penalty' for deviating too much on the issue in question, as compared to other issues, so *ceteris paribus* there is more confidence in judgements concerning issues with higher weights, under this confidence ranking. This is borne out by the following proposition.

[11] In particular, they do not automatically provide a distance on any space except $\Omega$.

**Proposition 2.** *Suppose, in (5), that* $w_L^L > w_R^L$ *and* $0.8w_B^L < w_L^L - w_R^L$*. Then, for every* $\epsilon \in (0, 0.9)$*, there exists* $o \in O$ *with* $c^L(o) \subseteq \mathcal{L}_\epsilon$ *but* $c^L(o) \nsubseteq \mathcal{R}_\epsilon$ *(where* $\mathcal{L}_\epsilon$*,* $\mathcal{R}_\epsilon$ *are as defined in Example 3.1).*

Whenever $w_B^L$ is not too large, if $w_L^L > w_R^L$, then any judgement $\mathcal{L}_\epsilon$ that the probability of $L$ is higher than a deviation $\epsilon$ below its best-guess probability 0.9 is held with more confidence than a judgement about $R$ that involves the same divergence $\epsilon$ from its best-guess probability 0.1 ($\mathcal{R}_\epsilon$). Figure 3a illustrates the intuition: when $w_L^L > w_R^L$, the sets in the confidence ranking have an 'elliptical' shape which is thinner along the $L$ dimension, hence translating higher confidence in judgements on this issue. So $w_L^L > w_R^L$ reflects higher confidence *ceteris paribus* in judgements about the labour market as compared to the real estate sector, and would be a natural assumption for Laura's confidence ranking, given her expertise. The clause concerning $w_B^L$ is related to the constraints that a given value of $p(B)$ places on the possible values of $p(L)$ and $p(R)$, as will be discussed below.

If a confidence ranking of the form (5) with $w_L^L > w_R^L$ is a natural representation for Laura, the corresponding confidence ranking for Ray is centred on $p^R$ as specified in Table 1,[12] with weights $\mathbf{w^R} = (w_L^R, w_R^R, w_B^R)$ where $w_R^R > w_L^R$, translating his relative expertise in real estate.

## 3.4 Confidence aggregation with within-person expertise diversity

Armed with confidence rankings reflecting cross-issue differences in expertise, we now consider aggregation of such rankings. The following Proposition characterises the centre of the confidence ranking obtained by confidence aggregation with the average confidence-level aggregator.

**Proposition 3.** *Consider a profile of confidence rankings of the form (3), with classical distance d, centre* $p^i$ *and vector of positive real-valued weights* $\mathbf{w^i}$ *for each* $i = 1, \ldots, n$*. Then the centre of the consensus-preserving confidence aggregation under the average confidence-level aggregator is:*

[12] I.e. $p^R$ such that $p^R(\omega_{LR}) = 0.09,\ p^R(\omega_L) = 0.01,\ p^R(\omega_R) = 0.81, p^R(\omega_N) = 0.09.$

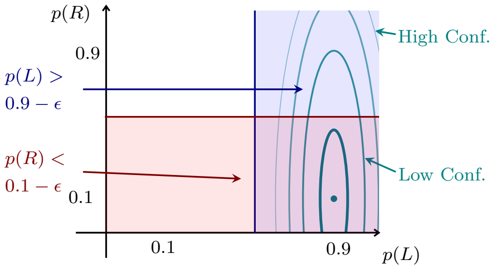


(a) Illustration of Proposition 2.
**Note:** As in Figure 2a, the blue area represents the probability judgement, $\mathcal{L}_\epsilon$, that $p(L)$ is within $\epsilon$ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement, $\mathcal{R}_\epsilon$, that $p(R)$ is within $\epsilon$ of $p^L(R) = 0.1$. The confidence in these judgements corresponds to the largest elliptical set contained in each area (Section 2.1): it is higher for the judgement concerning $L$.

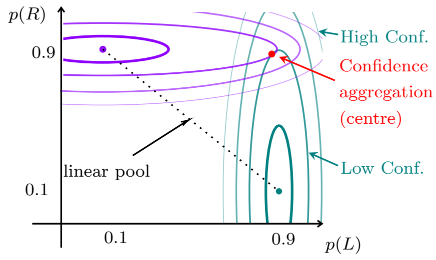


(b) Illustration of expertise-sensitive aggregation (Example 3.4).
**Note:** The red point is the centre of the result of confidence aggregation applied to the two confidence rankings, which coincides with expertise-sensitive pooling (Definition 5). The aggregate probability of $L$ is closer to Laura's judgement ($p^L(L) = 0.9$), and similarly for $R$. The dotted line is the set of points obtained by linear pooling (with different weights).

Figure 3: Confidence rankings generated as in Eq. (5).
**Note:** Each graph shows the space of pairs of probability values $(p(L), p(R))$ for the Labour and Real Estate events ($L$ and $R$; Example 3.1). The areas (sets of probability values) enclosed by the green ellipses represent the projection into this space of the $\mathbf{w^L}$ Euclidean-confidence ranking generated by $p^L$—i.e. Eq. (5)—with $w^L_L > w^L_R$ and $w^L_B$ low, representing Laura's confidence in beliefs. Larger, lighter ellipses correspond to higher confidence levels. The purple ellipses represent the $\mathbf{w^R}$ Euclidean-confidence ranking generated by $p^R$ (representing Ray), with $w^R_L < w^R_R$ and $w^R_B$ low.

$$\arg\min_{p \in \Delta} \sum_{i=1}^{n} \sum_{j=1}^{m} w^i_j d(p|_{\mathcal{P}_j}, p^i|_{\mathcal{P}_j}) \tag{6}$$

As we now show on the running example, this aggregation naturally incorporates within-person cross-issue expertise diversity.

**Example 3.4.** Consider, for Laura and Ray, the confidence rankings defined in Example 3.3 with $w^L_L > w^L_R$ and $w^R_L < w^R_R$. As discussed above, these rankings faithfully reflect Laura's higher expertise on the labour issue as compared to the real estate one, and similarly for Ray. Moreover, the example stipulates that Laura has more expertise in the labour market than Ray; in the light of the analysis of confidence rankings of form in Eq. (5), this suggests that $w^L_L > w^R_L$. Similarly, given Ray's higher specialisation in the real estate sector, $w^R_R > w^L_R$. Note that $\frac{w^L_B + w^R_B}{w^L_L + w^R_L}$ reflects the ratio of the overall confidence in the probability judgements on $B$ (across both agents) to the overall confidence in judgements concerning $L$, and similarly for $\frac{w^L_B + w^R_B}{w^L_R + w^R_R}$.

The following Proposition illustrates some properties of confidence aggregation with such confidence rankings.

**Proposition 4.** *Suppose that $c^L$ is specified as in Eq. (5) and similarly for $c^R$, and that $w^L_L > w^R_L$ and $w^R_R > w^L_R$. Then the centre in (6) contains a single probability measure, which we denote $p^{rep}$. Moreover:*

*i. If $\frac{w^R_L}{w^L_L} = \frac{w^L_R}{w^R_R}$ and $\frac{w^L_B + w^R_B}{w^L_L + w^R_L}, \frac{w^L_B + w^R_B}{w^L_R + w^R_R} \leqslant 1$ then $|p^L(L) - p^{rep}(L)| < |p^{rep}(L) - p^R(L)|$ and $|p^R(R) - p^{rep}(R)| < |p^{rep}(R) - p^L(R)|$;*

*ii. As $\frac{w^L_B + w^R_B}{w^L_L + w^R_L} \to 0$ and $\frac{w^L_B + w^R_B}{w^L_R + w^R_R} \to 0$:*

$$
\begin{aligned}
p^{rep}(L) \to & \frac{w^L_L}{w^L_L + w^R_L} p^L(L) + \frac{w^R_L}{w^L_L + w^R_L} p^R(L) \\
p^{rep}(R) \to & \frac{w^L_R}{w^L_R + w^R_R} p^L(R) + \frac{w^R_R}{w^L_R + w^R_R} p^R(R) \\
p^{rep}(B) \to & \begin{cases} 0.09 & \textit{if } p^{rep}(L) + p^{rep}(R) - 1 \leqslant 0.09 \\ p^{rep}(L) + p^{rep}(R) - 1 & \textit{otherwise} \end{cases}
\end{aligned}
$$

Recall from the Introduction that, under linear pooling, if the aggregate probability is closer to Laura's judgement than Ray's on one of the issues, then the same holds for the other issue. Proposition 4 shows that this is typically not true under confidence aggregation when confidence rankings reflect expertise. Part i. focuses on the case of symmetric expertise—Ray's expertise shortfall as compared to Laura on labour ($\frac{w^R_L}{w^L_L}$) is the same as her expertise shortfall compared to him on real estate ($\frac{w^L_R}{w^R_R}$). When the overall confidence in the judgements about the issues $L$ and $R$ outweigh the confidence in the judgement concerning the joint issue $B$, the resulting probability is closer to that of the individual with more expertise—both for the labour issue and for real estate. Given that, in the examples, Laura's expertise concerns $L$ but not specifically $R$ or $B$, and similarly for Ray, it seems reasonable that the overall confidence will be higher for the issues $L$ and $R$ than for $B$.[13]

Indeed, in the example, the overall confidence in the individual issues appears to be much higher than in the joint issue. Part ii. considers the limit as this confidence gap increases. It shows that the centre probability for $L$, $p^{rep}(L)$, tends to the weighted average of Laura's and Ray's judgements on $L$, where the weights are those in the generation of the confidence rankings that correspond to the issue $L$. Since Laura has more expertise than Ray on the labour market, $w_L^L > w_L^R$, so the probability for $L$ will be closer to Laura's; and similarly for $p^{rep}(R)$ and Ray's judgement on real estate. So again, this part of the result illustrates a typical situation (with potentially asymmetric expertise) where, under confidence aggregation, the resulting probability is closer to that of the relevant expert on both issues.

Figure 3b provides a visual illustration: the centre under confidence aggregation belongs to sets with confidence levels that are not too high on either ranking, and this picks out probability measures that are close to both Laura's probability on $L$ and Ray's on $R$. Confidence aggregation applied to these confidence rankings, which reflect within-person cross-issue expertise differences, thus follows each individual more closely on their area of expertise. It fairs better on this score than standard pooling rules.

Given that the centres of the confidence rankings generated as in (5) are probability measures assigning probability 0.09 to $B$, the centre of the resulting ranking will stick as close to this value as possible. If a probability 0.09 for $B$ is consistent with those of $L$ and $R$, then $p^{rep}(B)$ takes this value. If not, as will typically be the case, then $p^{rep}(B)$ takes the value closest to 0.09 that satisfies the constraints, i.e. $p^{rep}(L) + p^{rep}(R) - 1$. Since this is typically not 0.09,[14] Proposition 4 also demonstrates that the confidence aggregation rule does not respect spurious unanimities (Section 1).[15]

Examples 3.3-3.4 illustrate a simple way in which the confidence ap-

[13]This result is provided as a simple illustration. Stronger but more complicated versions are derivable from the calculations reported in Appendix A.

[14]E.g. when $w_L^L = w_R^R = 0.75$ and $w_R^L = w_L^R = 0.25$ in part ii., $p^{rep}(B) = 0.4$.

[15]Note that while this example involves 'primitive' views on the joint issue $B$, an alternative would be to work with opinions on the relationship between $L$ and $R$. Appendix A.2 shows how the approach set out above extends naturally to such cases.

proach can capture within-person expertise diversity, showing that confidence aggregation faithfully reflects these expertise differences in the resulting beliefs. It thus resolves the within-person cross-issue expertise diversity challenge. Moreover, it discounts agreements in the individuals' judgements about $B$ when there is little comparative confidence in them, thereby tackling the spurious unanimity challenge as well.

This suggests that, when applied on individuals' actual confidence rankings reflecting their acknowledged differences in expertise across issues, confidence aggregation will yield group beliefs that incorporate within-person cross-issue expertise differences. As argued in Section 3.3, confidence rankings such as those in Definition 4 are often a closer representation of individuals' beliefs than those underlying standard pooling rules; Propositions 3 and 4 show that, applied to these rankings, confidence aggregation respects expertise differences across issues.

Within-person cross-issue expertise diversity is equally relevant under the representative's-confidence interpretation, whereby confidence rankings reflect the representative's confidence in individuals' judgements (Section 3.1). In the example, the representative is justified in being more confident in Laura's judgement about labour than Ray's—and inversely for real estate. It is thus more reasonable to use expertise-sensitive generated confidence rankings (Definition 4) rather than classical distance-based ones (Proposition 1). As we now show, this generates a new pooling rule.

### 3.5 Expertise-sensitive pooling

Confidence aggregation applied on confidence rankings generated by classical distances recoups standard pooling rules (Proposition 1); we now consider what we obtain when using $\mathbf{w^i}$ $d$-confidence rankings instead (Definition 4). To this end, consider the correspondence yielding the centre in the result of confidence aggregation applied to $\mathbf{w^i}$ $d$-confidence rankings.

**Definition 5.** Let $\mathcal{P}_1, \ldots, \mathcal{P}_m$ be a set of partitions, and $d$ a classical distance. The correspondence $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m} : \Delta^n \rightrightarrows \Delta$ is defined by

$$F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}(p^1, \ldots, p^n) = \arg\min_{p\in\Delta} \sum_{i=1}^{n} \sum_{j=1}^{m} w^i_j d(p|_{\mathcal{P}_j}, p^i|_{\mathcal{P}_j}) \tag{7}$$

where $\mathbf{w^i} = (w_1^i, \ldots, w_m^i)$ is a tuple of vectors of positive real-valued weights, one for each individual.

As yet, $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ is not a well-defined probability aggregation rule. In particular, since the optimisation problem may have multiple solutions, $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ may yield a set of probability measures rather than a unique one. However, under some natural conditions, $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ is a well-defined (i.e. single-valued) probability aggregation rule.

**Proposition 5.** *Let $\mathcal{P}_1, \ldots, \mathcal{P}_m$ be a rich set of partitions, and $d$ a convex distance. Then $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ is a well-defined pooling rule, i.e. a function from $\Delta^n$ to $\Delta$. Moreover, the same holds for a rich set partitions and $d$ being the reverse relative entropy distance.*

Confidence aggregation thus generates a new pooling rule, which we call *expertise-sensitive pooling*. Given the remarks in Section 2.1, versions exist for all the classical statistical distances in Table 2. Since this is basically the rule used in Example 3.4, all the conclusions there—including natural respect for within-person expertise diversity—apply equally for this pooling rule. Indeed, as discussed in Section 6, the limit expressions concerning $L$ and $R$ in Proposition 4 are reminiscent of early suggestions in the pooling literature; unlike those suggestions, however, expertise-sensitive pooling is well-defined.

As an illustration of the technical simplicity of the confidence approach, we set out the intuition behind the first clause of Proposition 5, for interested readers. It relies on the observation that the optimisation problem defining $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ can typically be reduced to a recognisable, tractable form. More specifically, for a sequence of partitions $\mathcal{P}_1, \ldots \mathcal{P}_m$, let $P_{\mathcal{P}_1,\ldots\mathcal{P}_m} = \{(p|_{\mathcal{P}_1}, \ldots, p|_{\mathcal{P}_m}) \in \prod_{k=1}^m \Delta(\mathcal{P}_k) : p \in \Delta\}$, i.e. the set of sequences of probability measures on the partitions derived from some probability measure on $\Omega$. Since projection is a linear map, $P_{\mathcal{P}_1,\ldots\mathcal{P}_m}$ is a convex set. In fact, it is typically defined by a collection of inequalities: for instance, in the case of our running example (Example 3.2), the set is defined by a system of three linear inequalities (see Appendix A.1 for details). The centre of the aggregate confidence ranking (7) can thus be equivalently characterised as the set of probability measures $p$ such that $(p|_{\mathcal{P}_1}, \ldots, p|_{\mathcal{P}_m})$ belongs to:

$$\underset{(p_1,\ldots,p_m)\in P_{\mathcal{P}_1,\ldots,\mathcal{P}_m}}{\arg\min} \sum_{i=1}^{n}\sum_{j=1}^{m} w_j^i d(p_j, p^i|_{\mathcal{P}_j}) \tag{8}$$

For many settings of the parameters, problems of this sort are well-known: for instance, for the **w** Euclidean-confidence rankings in Example 3.3, this is a quadratic optimisation problem over a convex set (Appendix A.1). In general, whenever $d$ is convex, (8) is a minimisation of a strictly convex lower semicontinuous function on a convex, compact set, so there is a unique minimum. So whenever $\mathcal{P}_1, \ldots, \mathcal{P}_m$ is rich, (8) defines a unique probability measure in $\Delta$, hence establishing the Proposition.

# 4 Characterising Confidence Aggregation

In this section we provide a preference-based axiomatisation of consensus-preserving confidence aggregation (Definition 3) in a single-profile variable-population setting. In terms of the general aggregation-for-subsequent-decision problem (Section 2), here we shall thus be concerned with the relationship between the representative's preferences and the individuals' input. We begin by setting out the decision framework and preference representation.

## 4.1 Preliminaries: Representative's and consensus preferences

**Setup** Let $\mathcal{X}$, the set of *consequences*, be a convex subset of a vector space. $\mathcal{A}$ is the set of *acts*: simple measurable functions from states $\Omega$ to consequences $\mathcal{X}$. Mixtures of acts are defined pointwise as standard: for any $f, g \in \mathcal{A}$ and $\alpha \in [0,1]$, the $\alpha$-mixture of $f$ and $g$, which we denote with $f_\alpha g$, is defined by $f_\alpha g(\omega) = \alpha f(\omega) + (1-\alpha)g(\omega)$ for all $\omega \in \Omega$. We denote by $\mathcal{A}^c$ the set of a constant acts, i.e. acts yielding the same consequence in all states.

We use $\succ$ (perhaps with superscripts) to denote a strict preference relation on $\mathcal{A}$. $\succeq$ is the derived weak preference, defined as $f \succeq g$ if and only if $g \not\succ f$. Preferences $\succ$ *contradict* $\succ'$ if there exists $f, g \in \mathcal{A}$ with $f \succ g$ and $f \prec' g$. A preference relation $\succ$ is *contradictory* if there exists

$f, g \in \mathcal{A}$ with $f \succ g$ and $f \prec g$. As standard, a functional $V : \mathcal{A} \to \mathbb{R}$ is said to represent $\succ$ if, for all acts $f, g \in \mathcal{A}$, $f \succ g$ if and only if $V(f) > V(g)$.

**Representative's preferences and decision models** It is well known that ambiguity-sensitive, complete preferences can be connected to underlying incomplete preferences (Ghirardato et al., 2004; Gilboa et al., 2010). The bridge between the two concerns ambiguity attitudes—a taste factor, which, in the context of the aggregation-for-subsequent-decision problem, typically lies under the competence of the decision maker, not the experts. Hence, whilst the decision maker's (weak) preferences are complete, it is natural to allow the representative's relevant preferences to be incomplete. We thus follow Danan et al. (2016) in studying aggregation in the context of incomplete preferences. As noted in Section 2.1, the representation of confidence in beliefs used here is compatible with a wide range of models of decision under uncertainty. It supports a general incomplete preference representation (Hill, 2016) according to which, for all acts $f, g \in \mathcal{A}$, $f \succ g$ if and only if:

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \qquad \text{for all } p \in c(\max\{D(f), D(g)\}). \tag{9}$$

where $\mathbb{E}_p$ is the expectation with respect to a probability measure $p \in \Delta$,[16] $u : \mathcal{X} \to \mathbb{R}$ is a non-constant affine utility function, $c$ is a closed convex confidence ranking and $D : \mathcal{A} \to O$ is such that $D(x) = \min O$ for all $x \in \mathcal{A}^c$ and it satisfies the following *richness* condition: for every $(f, g) \in \mathcal{A} \times \mathcal{A} \backslash \mathcal{A}^c \times \mathcal{A}^c$ and $o \in D(\mathcal{A} \backslash \mathcal{A}^c)$, there exists $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $\max\{D(f_\alpha h), D(g_\alpha h)\} = o$. Under this model, the confidence level used to make a decision may depend on the stakes involved in it; the function $D$, called the *cautiousness coefficient*, specifies the confidence level the decision maker considers relevant for evaluating each act, and hence for taking each decision. As shown in Hill (2013, 2016), it captures the decision maker's attitudes to choosing on the basis of limited confidence—a taste factor.

If (9) holds, we say that $(c, D, u)$ represents $\succ$. The representing $u$ is unique up to positive affine transformation, $c(O)$ is unique up to convex closure and $c(\max\{D(\bullet), D(\bullet)\})$ is unique (Hill, 2016).

[16] I.e. for any $\phi : \Omega \to \mathbb{R}$, $\mathbb{E}_p \phi = \int \phi(\omega) dp(\omega)$.

Note finally that (9) is a generalisation of Bewley (1986) preferences, according to which, for all acts $f, g \in \mathcal{A}$, $f \succ g$ if and only if:

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \qquad \text{for all } p \in \mathcal{C}. \tag{10}$$

where $u$ is as above and $\mathcal{C}$ is a closed, convex subset of $\Delta$. For every non-contradictory preference $\succ$, define the *Bewley min-extension* $\widehat{\succ}$ as the relation containing every comparison that holds in every Bewley preference extending $\succ$:

$$\widehat{\succ} = \bigcap_{\substack{\succ' : \succ \subseteq \succ' \\ \succ' \text{ represented by (10)}}} \succ' .$$

**General aggregation problem** Recall that $I$ is a finite set of individuals. Each individual $i$ provides fixed input $\mathcal{I}^i$. As in single-profile contexts that are common in the related literature (e.g. Mongin, 1995; Danan et al., 2016), the sequence of inputs $(\mathcal{I}^i)_{i \in I}$ will be fixed throughout. However, we consider variable-population single-profile aggregation: the representative has access to the inputs from a set $J \subseteq I$ of individuals, which may vary. For each non-empty $J \subseteq I$, $\succ^J$ is the representative's preference formed on the basis of inputs $(\mathcal{I}^i)_{i \in J}$ from the individuals in $J$. As a point of notation, for singleton $J = \{i\}$, we use $\succ^i$ in place of $\succ^{\{i\}}$.

As noted in Section 2, we make no assumption on the format of inputs; accordingly, this framework is rich enough to cover several standard setups in the literature. For instance, each $\mathcal{I}^i$ could be a probability measure $p^i \in \Delta$, as in the case of probability aggregation or model misspecification. In this case, the preferences $\succ^i$ would be based on the confidence rankings generated by these probabilities, as under the representative's-confidence interpretation (Section 3.1). Alternatively, each $\mathcal{I}^i$ could be a preference satisfying (9). Whenever the representative's preferences $\succ^i$ coincide with the input preferences $\mathcal{I}^i$, this reduces to a classic preference aggregation problem (Crès et al., 2011; Danan et al., 2016) for confidence preferences.

We assume that the representative's preferences are represented according to the incomplete preference model (9), with the same tastes no matter the inputs received. Since this model has two parameters representing tastes—the utility function $u$ and the cautiousness coefficient $D$—this is expressed by the following assumption.[17]

**Assumption 1.** *For each non-empty $J \subseteq I$, $\succ^J$ can be represented according to (9) with the same $D$ and $u$ for all $J$.*

When the inputs are confidence preferences, this corresponds to the standard assumption in the aggregation literature that all individuals and the group share the same tastes (e.g. Crès et al., 2011).

Henceforth, we fix representations $(c^J, D, u)$ of $\succ^J$, for non-empty $J \subseteq I$. As for preferences, we use $c^i$ as shorthand for $c^{\{i\}}$.

**Stakes** A central idea behind representation (9) and related models is that the more important a decision is, the greater the confidence required in the beliefs one relies on to make it (Hill, 2013, 2016, 2019b; Bradley, 2017a). For instance, (9) represents agents for which determinate preferences held at low stakes—where less confidence is required—may become indeterminate at higher stakes. In the light of this, the use of higher-confidence beliefs—i.e. $\max\{D(f), D(g)\} > \max\{D(f'), D(g')\}$—indicates that the agent considers the choice between $f$ and $g$ to be more important than the choice between $f'$ and $g'$: it involves higher stakes. In the context of (9), this can be formalised by a surjective function $\sigma : \mathcal{A} \times \mathcal{A} \to \mathcal{S} \subseteq \mathbb{R}_{\geqslant 0}$, assigning to each binary choice the stakes involved in it. Hill (2016) contains several examples of such (real-valued) notions of stakes.[18] For $(f, g) \in \mathcal{A} \times \mathcal{A} \backslash \mathcal{A}^c \times \mathcal{A}^c$ and $s \in \mathcal{S}$, we say that $(f, g)$ *has stakes* $s$ if $\sigma(f, g) = s$. Assumption 1 guarantees that the representative's preferences at different $J$ are consistent with a single stakes function $\sigma$, in the sense that $D$ and $\sigma$ are linked by a monotone transformation.[19]

[17]Given the uniqueness of the representation, this is equivalent to the uniqueness of $u$ and $D$ up to appropriate transformations. Indeed, whilst stated on the models for ease, this Assumption can be reformulated in behavioural terms, drawing, for instance, on the choice-based foundations for the weak preference version of (9) provided by Hill (2016), which can be adapted to strict preference using Bewley, 1986; Karni, 2011.

[18]Hill (2016) also discusses the relationship between the stakes involved in an act and the stakes in a (binary) choice: in particular, whilst, for simplicity, we adopt one possible relationship in (9), Hill (2016) contains examples of others.

[19]More precisely, for all $f, g \in \mathcal{A}$, $\max\{D(f), D(g)\} = \zeta \circ \sigma(f, g)$ for some strictly increasing continuous real-valued function $\zeta$.

Given a preference relation $\succ$ represented according to (9) and a stakes level $s \in \mathcal{S}$, the derived relation $\succ_s$ is the relation extending $\succ|_{\mathcal{A}^c \times \mathcal{A}^c}$ such that: for all $(f, g) \in \mathcal{A} \times \mathcal{A} \backslash \mathcal{A}^c \times \mathcal{A}^c$, $f \succ_s g$ if and only if there exists $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $(f_\alpha h, g_\alpha h)$ has stakes $s$ and $f_\alpha h \succ g_\alpha h$.[20] As discussed in Hill (2013, 2016), $f \succ_s g$ essentially says that, if the acts were evaluated 'as if' the decision involved stakes $s$, then $f$ would be preferred. For example, consider two choices. One is between the bet $f_{2028}$ on the Democrat candidate winning the 2028 US President election, yielding \$1 million if you win and a loss of \$1 million if not, and nothing, $g_{2028}$. The other choice is between a similar bet $f_{2032}$ on the 2032 election, with stakes (winnings and losses) a million times less in utility terms, and no utility change, $g_{2032}$. An agent with beliefs that are more precise and slightly more favorable for the 2028 bet might nevertheless choose the bet in the 2032 choice but have indeterminate preferences in the 2028 one because of the difference in stakes: with lower stakes, he can rely on low-confidence beliefs when comparing $f_{2032}$ and $g_{2032}$, but not for the choice between $f_{2028}$ and $g_{2028}$. However, if the 2028 choice were evaluated at the low stakes level, say $s$, then $f_{2028}$ would typically be chosen over $g_{2028}$: i.e. $f_{2028} \succ_s g_{2028}$. When $f \succ_s g$, we say that $f$ *is preferred to* $g$ *at stakes level* $s$, and we call $\succ_s$ *the preferences at stakes level* $s$.

**Consensus preferences** A *profile of stakes levels* is a variable-population tuple of stakes levels $\mathbf{s} = (s_i)_{i \in J} \in \bigcup_{J \subseteq I} \mathcal{S}^{|J|}$, where $s_i$ is understood to be the stakes level for the $i^{\text{th}}$ individual under $\mathbf{s}$. Each profile $\mathbf{s}$ yields a profile of preferences: those at the relevant stakes level for each individual involved. There is a consensus if these preferences do not contradict each other; hence the following definition.

**Definition 6.** For a profile of stakes levels $\mathbf{s} = (s_i)_{i \in J} \in \bigcup_{J \subseteq I} \mathcal{S}^{|J|}$, $\mathbf{s}$ *exhibits consensus* when $\bigcup_{i \in J} \succ^i_{s_i}$ is not contradictory; it does not exhibit consensus otherwise. If $\mathbf{s}$ exhibits consensus, the *consensus preference* $\succ_{\mathbf{s}}$ on $\mathcal{A}$ is defined by $\succ_{\mathbf{s}} = \widehat{\bigcup_{i \in J} \succ^i_{s_i}}$. $\succ^J$ *respects the consensus* $\succ_{\mathbf{s}}$ *at stakes level* $s$ if $\mathbf{s}$ exhibits consensus and $\succ^J_s \subseteq \succ_{\mathbf{s}}$.

Consensus-preserving confidence aggregation aims to respect consensus

[20]This is well-defined by the richness of $D$.

borne of compromise (Section 2.2). The union of all the (determinate) preferences at the stakes levels specified by $\mathbf{s}$ naturally reflects a consensus whenever none of these preferences contradict each other. The corresponding consensus preference $\succ_{\mathbf{s}}$ assembles these preferences and their 'objectively rational' entailments, i.e. those comparisons that hold under every Bewley extension (Gilboa et al., 2010). This coherent consensus position is the product of compromise: it results from each individual $i \in J$ only putting their preferences at stakes level $s_i$ 'on the table', and leaving other preferences aside. The representative's preference $\succ^J$ respects the consensus $\succ_{\mathbf{s}}$ at stakes level $s$ if it does not go beyond it: every preference in $\succ^J_s$ appears in the consensus, though some consensus preferences may be left open by the representative at that stakes level.

## 4.2 Confidence aggregation and Pareto

**Axioms** The preference-based characterisation of confidence aggregation relies on one main axiom. To introduce it, first consider the Pareto axiom behind linear pooling in a sufficiently rich, single-profile aggregation context (Mongin, 1995). The strict preference version (in our setup) is as follows.

**Axiom** (Standard Pareto)**.** *For all acts $f, g \in \mathcal{A}$ and non-empty $J \subseteq I$, if $f \succ^i g$ for all $i \in J$, then $f \succ^J g$.*

This principle encodes preservation of unanimity across individual preferences at the level of isolated comparisons between pairs of acts: if everyone holds a certain preference between two acts, then it should be adopted by the group's representative. As noted in the Introduction, this leads to challenges related to spurious unanimity. Consider the following variant.

**Axiom** (Consensus-wise Pareto)**.** *For all stakes levels $s \in \mathcal{S}$ and non-empty $J \subseteq I$ such that $\succ^J$ respects some consensus $\succ_{\mathbf{s}'}$ at $s$, and all acts $f, g \in \mathcal{A}$, if $f \succ_{\mathbf{s}} g$ for all $\mathbf{s}$ for which $\succ^J$ respects the consensus at $s$, then $f \succ^J_s g$.*

Rather than asking the representative to adopt a preference if every individual holds it, Consensus-wise Pareto looks across all relevant consensuses formed from the individuals' preferences. If the preference holds in all consensuses respected at a given stakes level, then the representative adopts that preference at those stakes. Consensus-wise Pareto thus retains the motivation that group-level preferences should respect consensuses among its members, but, unlike Standard Pareto, it does not conceive of consensus as simple unanimity in the preference comparison between two acts. Rather, it respects the richer consensus preferences $\succ_{\mathbf{s}}$, understood as coherent overall positions acceptable to all individuals, insofar as they combine their preferences whilst compromising by leaving lower-stakes ones 'off the table'.

Note that more consensuses are respected at higher stakes levels than at lower ones, so fewer preferences hold in all such consensuses. This principle thus applies to fewer preferences at higher stakes levels, in line with the expectation that fewer preferences are held with higher confidence.

Whilst, logically, neither Standard Pareto nor Consensus-wise Pareto imply the other, Proposition 1 shows that linear pooling can be recovered as a special case of confidence aggregation. In this sense, the latter condition is more general.

Our characterisation requires two auxiliary axioms.

**Axiom** (Consensus-based preferences)**.** *For all stakes levels $s \in \mathcal{S}$, acts $f, g \in \mathcal{A}$ and non-empty $J \subseteq I$, if $f \not\succ^J_{s'} g$ for every stakes level $s'$ such that some consensus $\succ_{\mathbf{s}}$ is respected at $s'$, then $f \not\succ^J_s g$.*

**Axiom** (Non-degeneracy)**.** *There exists a tuple of stakes levels $\mathbf{s} = (s_i)_{i \in I}$ exhibiting consensus.*

Aggregate preferences should come from individuals' preferences. Under confidence aggregation, the latter constrain the former principally in the context of consensuses, and hence at stakes levels where some consensus is respected. Consensus-based preferences states that all of the representative's preferences are determined by those formed on the basis of consensuses: in particular, any preferences at a stakes level where no consensus is respected must already be present at a level where some are. Non-degeneracy states that, if individuals leave sufficiently many preferences aside by restricting to high enough stakes levels, they can come to a consensus.

**Result** We have the following characterisation result.

**Theorem 1.** *Let $\{\succ^J\}_{J\subseteq I}$ satisfy Assumption 1. They satisfy Consensus-wise Pareto, Consensus-based preferences and Non-degeneracy if and only if, for each $J$ with $|J| \geqslant 2$, $c^J$ is a consensus-preserving confidence aggregation of $(c^i)_{i\in J}$ up to convex closure.*

*Moreover, for each $J \subseteq I$ with $|J| \geqslant 2$, there is a unique minimal confidence-level aggregator $\otimes$ under which $c^J$ is a consensus-preserving confidence aggregation: that is, for all $\otimes'$ such that $c^J$ is a consensus-preserving confidence aggregation of $(c^i)_{i\in J}$ under $\otimes'$ up to convex closure, $\otimes'(\mathbf{o}) \geqslant \otimes(\mathbf{o})$ for all $\mathbf{o}$ such that $\bigcap_{i\in J} c^i(o_i) \neq \varnothing$.*

So the central axiom characterising confidence aggregation is Consensus-wise Pareto, which is just a reformulation of the standard Pareto condition to apply to a more refined notion of consensus rather than unanimity in individual preferences at the level of isolated comparisons. Indeed, Consensus-based preferences can be dropped whenever the representative's confidence ranking is pointed. More generally, without it, the representative's confidence ranking is always that obtained by a confidence aggregation, except at confidence levels at the bottom of the ranking.

No particular confidence-level aggregator is assumed in this result; rather, the appropriate aggregator is fixed endogenously by preferences. Moreover, there is a unique minimal one: that is, one which always takes the lowest value across all aggregators representing the preferences. Relatedly, no particular confidence-level aggregator is imposed in Theorem 1. Extensions involving specific confidence-level aggregators or relationships between the aggregators for different $J$ lie largely beyond the scope of this paper. To illustrate some possibilities, Appendix B contains supplementary axioms that characterise some confidence-level aggregators related to those in Table 3 (Section 2.2).

# 5 Decision making, model misspecification and confidence

Building on the previous developments concerning aggregation, we now bring in decision. This will permit discussion of attitudes to uncertainty, and notably the relationship with recent model misspecification decision rules. In terms of the general aggregation-for-subsequent-decision problem, this requires considering the relationship between the representative's preferences, as studied in the previous section, and the decision maker's. We begin by characterising this link, to obtain a representation of the decision maker's preferences combining the aggregation and decision dimensions.

## 5.1 Decision maker's preferences

Let $\succ^{rep}$ denote the representative's preferences resulting from the aggregation stage studied in Section 4, and $\succ^{dm}$ the decision maker's. We adopt standard 'rationality' assumptions on the decision maker's preferences.

**Assumption 2** (Decison maker's preferences). $\succeq^{dm}$ *are complete, transitive and continuous.*

Consider the following axioms relating the representative's and decision maker's preferences.

**Axiom** (C-Consistency). *For all* $f \in \mathcal{A}$, $x \in \mathcal{A}^c$,

$$f \succ^{rep} x \Rightarrow f \succ^{dm} x$$

**Axiom** (Caution). *For all* $f \in \mathcal{A}$, $x \in \mathcal{A}^c$

$$f \not\succ^{rep} x \Rightarrow x \succeq^{dm} f$$

These are strict-preference versions of standard axioms relating incomplete and complete preferences (Gilboa et al., 2010; Cerreia–Vioglio et al., 2025), albeit under different interpretations. C-Consistency is the reformulation of the Consistency axiom from Gilboa et al. (2010) in terms of strict preferences and its weakening to comparisons with constant acts. In the current context, it says that, in all such comparisons, the decision maker adopts all of the representative's preferences. Caution is the strict-preference version of Gilboa et al.'s Caution. In this context, it says that the decision maker does not adopt any strict preference for an act over a constant act whenever the representative does not do so.

**Proposition 6.** *Under Assumptions 1 and 2, the following are equivalent:*

*i.* $\succ^{rep}$ *and* $\succ^{dm}$ *jointly satisfy C-Consistency and Caution;*

*ii.* $\succ^{dm}$ *is represented by*

$$\min_{p \in c^{rep}(D(f))} \mathbb{E}_p u(f) \tag{11}$$

*where* $(c^{rep}, D, u)$ *represent* $\succ^{rep}$ *according to* (9).

Representation (11) encompasses a range of popular models in the ambiguity literature. It is the ambiguity averse model in the confidence family (Hill, 2013, 2019b), and as such generalises Gilboa and Schmeidler's (1989) maxmin-EU model. Like maxmin-EU, decision criterion (11) evaluates acts by the worst-case expected utility over a set of priors; unlike it, the set used varies according to the act evaluated, with higher-stakes acts being evaluated using sets corresponding to higher confidence levels. Moreover, as shown in Appendix C, representation (11) generalises variational preferences (Maccheroni et al., 2006), which underpin the main model misspecification decision rules. In particular, every standardly-used variational preference—and all those used in the model misspecification literature—can be represented according to (11) (Proposition C.1).

## 5.2 Aggregation for subsequent decision

Combining Proposition 6, relating the decision maker's preferences to the representative's, with Theorem 1, on the formation of latter preferences from the individuals' input, yields a representation of the decision maker's preferences in terms of the aggregate confidence ranking.

**Corollary 1.** *Let* $\{\succ^J\}_{J \subseteq I}$ *satisfy Assumption 1 and Consensus-wise Pareto, Consensus-based preferences and Non-degeneracy, and let* $\succ^{dm}$ *satisfy Assumption 2 and, in conjunction with* $\succ^I$, *C-Consistency and Caution. Then* $\succ^{dm}$ *is represented by:*

$$\min_{\substack{p \in \Delta: \\ \otimes((\iota^i(p))_{i \in I}) \leqslant D(f)}} \mathbb{E}_p u(f) \tag{12}$$

*for some confidence-level aggregator* $\otimes$, *where, for each* $i \in I$, $(c^i, D, u)$

*represents* $\succ^i$.[21]

Representation (12) captures decision in the face of multiple information sources that proceeds by aggregation followed by ambiguity-sensitive decision. It admits the two interpretations discussed in Section 3.1. Under the individuals'-confidence interpretation, the decision maker chooses on the basis of the confidence aggregation of the individuals' beliefs. This result thus connects to the literature on preference aggregation under decision under uncertainty (e.g. Crès et al., 2011), providing a characterisation for the confidence approach. Under the representative's-confidence interpretation, the decision maker chooses on the basis of the representative's assessments, formed from the individuals' inputs and her evaluation of their relative expertise. This is relevant in cases where a decision must be taken on the basis of inputs in the form of probability measures. A relevant example is when they are provided by models.

### 5.3 Making decisions with models

As discussed in Section 3, when input comes from models, i.e. when $\mathcal{I}^i = p^i \in \Delta$ for each $i$, confidence rankings generated by the individual probability measure inputs under a distance $\rho$ and weights $w(i)$ can be used in the confidence aggregation rule. Plugging these into (12) yields the representation of the decision maker's preference by:

$$\min_{\substack{q\in\Delta: \\ \otimes_{i\in I} w(i)\rho(q,p^i)\leqslant D(f)}} \mathbb{E}_q u(f) \tag{13}$$

Aside from the utility and the weights, three parameters dictate the evaluation of a given act: $\rho$ (determining confidence in input probabilities); $\otimes$ (regulating aggregation); and the stakes of the act (determining the relevant confidence level $D(f)$).[22]

[21] Recall from Section 2.1 that $\iota^i$ is an equivalent formulation of $c^i$, and that confidence aggregation can be formulated in terms of $\iota^i$ (Section 2.2).

[22] As mentioned above and discussed in the literature (e.g. Hill, 2013, 2016, 2019b), $D(f)$ reflects the stakes involved in $f$ and the decision maker's attitude to choosing on the basis of limited confidence. For instance $D(f)$ takes the lowest value both when $f$ is a low stakes act and when the decision maker systemically uses this confidence level to choose, as a Bayesian would. In the current setup, which follows much of

This case of our aggregate-then-decide decision rule recovers many approaches to choosing on the basis of models in the literature.

**Theorem 2.** *Let the assumptions in Corollary 1 hold, and, for each $i \in I$, let $\mathcal{I}^i = p^i \in \Delta$.*

*Let $\succ^i$ be represented using a $w(i)$ $\rho$-confidence ranking generated by $p^i$, for individual-specific weights $w(i)$ (Definition 2). Then, for each of the first four rows of Table 5, if $\rho$ and $\otimes$ are as specified in the first two columns, then preferences $\succ^{dm}$ over acts with stakes specified in the third column are represented by the functional in the fourth column.*

*Moreover, if instead each $\succ^i$ is represented by a $\mathbf{w}(i)$ d-confidence ranking generated by $p^i$, for individual specific weight vectors $\mathbf{w}(i)$ and partitions $\mathcal{P}_1, \ldots \mathcal{P}_m$ (Definition 4), then, for each of the last three rows of Table 5, if $d$ and $\otimes$ are as specified in the first two columns, then preferences $\succ^{dm}$ over acts with stakes specified in the third column are represented by the functional in the fourth column.*

We now discuss these cases in turn.[23]

**Low stakes: Bayesian Model Averaging** In evaluating acts involving low stakes, sets of priors held with low confidence levels will be used; if these sets are singletons (i.e. the set in the subscript of the minimisation in (13) is a singleton for small enough $D(f)$), then (13) coincides with subjective expected utility. So, for confidence-level aggregators $\otimes$ and distances $\rho$ yielding the results of, say, linear pooling (Proposition 1), (13) coincides

the literature on preference-based axiomatisations for belief aggregation (Section 4) in essentially assuming that all agents share the same tastes (in this case, $u$ and $D$), the attitude aspect is fixed, whereas the stakes will vary with the choice at hand. However, relying on the belief-only development in Sections 2 and 3, one can extend to a model where the (confidence in) beliefs are aggregated according to the proposed rule, and the decision maker provides $u$ and $D$ that may differ from those of the representative. Since the interpretation of $D(f)$ is discussed in the literature and not the focus of the current paper, for conciseness we refer to the 'stakes involved in an act $f$' when discussing the factors determining the confidence level $D(f)$ used to evaluate it, leaving it understood that, in practice, these may also include attitudes. So a low stakes act is an $f$ for which $D(f)$ is the lowest confidence level; a high stakes act is one for which $D(f)$ is high.

[23] Although $I$ has been assumed to be finite here—and Table 5 is written for finite $I$—in the model uncertainty literature the set of models is often infinite. The points made here extend to infinite sets of models, under appropriate technical modifications.

with EU on a linear pool of the models, at low stakes. Confidence aggregation with the confidence decision model (12) thus subsumes Bayesian Model Averaging (Raftery et al., 1997; Steel, 2020), a popular approach to dealing with multiple models that involves the linear pool of the distributions provided by the various models, with weights determined by the posterior probabilities over them (Table 5, row 1). Changing the distance $\rho$ yields EU with geometric pooling (Table 5, row 2).

**Medium / High stakes: Model misspecification** At higher stakes, sets further up the confidence ranking will feature in (13). This is illustrated in rows 3 and 4, featuring representations that generalise Hansen and Sargent's (2001; 2008) constraint preferences, which are the special case of (13) with no aggregation (i.e. singleton $I$), relative entropy distance and identical $D(f)$ for all $f$. For instance, row 4 shows that the specification yielding geometric pooling (row 2) generates an 'average' robust control evaluation at higher stakes. Changing the confidence-level aggregator $\otimes$ yields a 'minimum' robust control evaluation (row 3), where the constraint set is determined by the weighted minimum relative entropy distance.

Hansen and Sargent (2001, 2008) also introduce multiplier preferences, which yield the same optima as their constraint preferences on various classes of decision problems. By similar reasoning, for sufficiently convex distances, confidence-level aggregators and sets of models, (13) will yield the same optima on these problems as:[24]

$$\min_{q \in \Delta} \left( \mathbb{E}_q u(f) + \lambda \otimes_{i \in I} w(i) \rho(q, p^i) \right) \quad (14)$$

for appropriate $\lambda$. (14) contains several representations in the recent literature on model misspecification; for instance, under equal weights, the settings of $\otimes$ and $\rho$ in row 3 yield the representation proposed by Hansen and Sargent (2022). It follows from the previous observations that this representation yields the same optima as that in row 3 of Table 5 on the decision problems studied by Hansen and Sargent (2008).

[24]Hansen and Sargent's proof for multiplier preferences relies on the Lagrange multiplier theorem (Luenberger, 1969), and hence on the convexity of the constraint ($R(q\|p)$ in their case) as a function of $q$. Sufficient convexity conditions ensure the theorem applies to (13) and (14).

| $\rho$ | $\otimes$ | Stakes | 'Overall' Decision rule |
|---|---|---|---|
| Reverse RE | Average | Low | $\mathbb{E}_{\sum_I \frac{w(i)}{\sum_{i\in I} w(i)} i} u(f)$<br>Bayesian Model Averaging & SEU (Steel, 2020) |
| RE | Average | Low | $\mathbb{E}_{\chi \prod_I i^{\frac{w(i)}{\sum_I w(i)}}} u(f)$<br>Geometric pooling & SEU (Dietrich, 2021)[25] |
| RE | Minimum | Medium / High | $\min_{\substack{q\in\Delta: \\ \min_{i\in I} w(i)R(q\Vert p^i)\leqslant\eta}} \mathbb{E}_q u(f)$<br>'Minimum' robust control<br>(Hansen and Sargent, 2022; Cerreia–Vioglio et al., 2025) [26] |
| RE | Average | Medium / High | $\min_{\substack{q\in\Delta: \\ \sum_I w(i)R(q\Vert p^i)\leqslant\eta}} \mathbb{E}_q u(f)$<br>'Average' robust control |
| Exp-sens. Reverse RE | Average | Low | $\mathbb{E}_{\arg\min_{q\in\Delta} \sum_I \sum_{j=1}^m w_j(i) R(p^i\vert_{\mathcal{P}_j}\Vert q\vert_{\mathcal{P}_j})} u(f)$<br>Expertise-sensitive pooling & SEU (Defn 5) |
| Exp-sens. RE | Minimum | Medium / High | $\min_{\substack{q\in\Delta: \\ \min_I \sum_{j=1}^m w_j(i)R(q\vert_{\mathcal{P}_j}\Vert p^i\vert_{\mathcal{P}_j})\leqslant\eta}} \mathbb{E}_q u(f)$<br>Expertise-sensitive minimum robust control |
| Exp-sens. RE | Average | Medium / High | $\min_{\substack{q\in\Delta: \\ \sum_I \sum_{j=1}^m w_j(i)R(q\vert_{\mathcal{P}_j}\Vert p^i\vert_{\mathcal{P}_j})\leqslant\eta}} \mathbb{E}_q u(f)$<br>Expertise-sensitive average robust control |

Table 5: Aggregation for subsequent decision: Special cases. See Tables 2 and 3 for definitions of values for $\rho$ and $\otimes$ respectively.

*Note:* RE stands for 'Relative Entropy' and Exp-sens. RE for 'Expertise-sensitive Relative Entropy', i.e. confidence rankings as in Definition 4 with $d$ the relative entropy classical distance; similarly for Reverse RE and Exp-sens. Reverse RE. Throughout $\eta \in \mathbb{R}_{>0}$. In the last three rows, $w_j(i) \in \Re_{\geqslant 0}$ is the weight assigned to model $i$ and issue $j$. Other notation is as in Sections 3, 4 and the text.

**Model expertise** Section 3 demonstrated that confidence aggregation can incorporate within-person cross-issue expertise diversity; this remains relevant when the 'experts' are models. In real applications, it is not un-

[25] As standard, the geometric pool involves a multiplicative constant, denoted $\chi$.

[26] Hansen and Sargent (2022) proposed the multiplier version of a special case of this model (i.e. (14) with $\otimes$ and $\rho$ as specified in the Table and $w(i) = 1$ for all $i \in I$); Cerreia–Vioglio et al. (2025, Proposition 6) axiomatise a version with general convex $\rho$, minimum $\otimes$ and convex, compact set of models. In such cases, the convexity assumptions needed to run Hansen and Sargent's (2008) argument hold (footnote 24), so the constraint preferences in the Table yield the same optima as the corresponding multiplier version in the relevant decision problems.

common for some models to be 'better' on certain issues and 'worse' on others. In climate science, say, one model could have a more detailed representation of cloud formation, whereas another is more accurate on elements of the biosphere: the former might thus be expected to do a better job in predicting rainfall, the latter in predicting ground-level temperature (Masson-Delmotte et al., 2021, Section 1.5.3). However, the distances underlying Bayesian Model Averaging and existing misspecification models (Table 5, rows 1–4) encode the assumption that all models have uniform expertise, or are equally 'good', across all issues (Section 3.3). Like standard pooling rules (Section 3.1), existing model-misspecification decision rules thus cannot cope with intra-model cross-issue expertise diversity.

Using the same expertise-sensitive confidence rankings (Definition 4) as in the derivation of the expertise-sensitive pooling rule (Definition 5) yields a expertise-sensitive generalisation of Bayesian Model Averaging (row 5) at low stakes. At higher stakes levels, it yields misspecification-sensitive decision rules that *can* accommodate differing degrees of expertise within models. Rows 6 and 7 provide two examples, obtained by replacing the distance generating existing misspecification models by the expertise-sensitive distance introduced in Section 3.3. For the reasons set out in Section 3, they constitute arguably more reasonable model misspecification decision rules in situations where models' performance may vary across issues.

Applying standard preference-completion axioms from the literature to the representative's preferences thus yields an aggregation-for-subsequent-decision framework combining confidence aggregation with an existing, general decision rule under uncertainty. It unifies a wide variety of accounts in the literature: in particular, with confidence rankings generated by classical statistical distances, it recovers both Bayesian and model misspecification approaches. Moreover, it naturally organises them along three dimensions: two aggregation ones—the generating distance reflecting assumptions about confidence (Section 3.2), and the only new parameter in the confidence aggregation rule (Section 2.2)—and one decision dimension—the appropriate confidence level for the decision, which ultimately reflects the decision maker's appraisal of the stakes involved and her attitudes to choosing on the basis of limited confidence (Hill, 2013, 2019b). The framework thus resolves the decision challenge by not only accommodating a range of approaches, but by situating the differences between them. For instance, as is clear from Table 5, Bayesian model averaging and the 'minimum' robust control rule (Hansen and Sargent, 2022) differ on both the aggregation and decision dimensions.

The separation of the aggregation and decision stages afforded by the confidence approach is comforting conceptually: the *epistemic* issue of the beliefs (and confidence) that can or should be formed on the basis of a set of models is different from the *pragmatic* question of their role in decision making. More importantly, it is frequently crucial in practice. As noted in the Introduction, the science-policy frontier is often structured along this distinction, with actors responsible for aggregating expert opinions about matters of fact, for use by a decision maker who supplies the tastes. By supporting this separation, the approach developed here naturally applies in such situations. Moreover, the separation facilitates the development of new decision rules, as in the case of the expertise-sensitive model misspecification rules in Table 5, which result from recognising the relevance of a challenge in the aggregation literature for model misspecification rules, and applying the solution developed in Section 3.

# 6 Discussion

As mentioned previously, the approach to the aggregation-for-subsequent-decision problem proposed here admits several interpretations. Under the representative's-confidence interpretation, the individuals' input can come in the form of probabilities; this connects our approach to the literature on probability aggregation taking probabilities as primitive, rather than working with preferences (Genest and Zidek, 1986). The within-person expertise diversity challenge was first raised in this literature, with some early contributions suggesting averaging with potentially different weights for each event (e.g. Bordley and Wolff, 1981). Such rules turned out not to be well-defined (e.g. McConway, 1981; Genest and Zidek, 1986). This, and in particular the apparent impossibility in capturing within-person expertise diversity, has long been argued to be a problem for linear pooling (e.g. French, 1985). The limit case in Example 3.4 (Section 3.4)—involving different weights for the labour and real estate events—shows that confidence aggregation can capture the intuition behind the early proposals. It does so whilst overcoming their limits: as testified by Proposition 5, the derived expertise-sensitive pooling rule is always well-defined.

The model misspecification literature also takes probabilities as primitives, with recent work focussing on cases of multiple probability measures (Hansen and Sargent, 2022; Cerreia–Vioglio et al., 2025). Section 5.3 shows that some of these decision rules can essentially be recovered as special cases of representation (13). However, unlike model misspecification models which bake everything into a single decision rule, the confidence framework separates the aggregation and decision dimensions. Beyond its practical relevance and support for the development of new approaches (Section 5.3), this separation can bring to light hitherto unrecognised relationships between models, as in the example of the affinities between the 'average' robust control and geometric pooling (Table 5, rows 2 & 4). On the decision front, extensions of representation (13) involving different functionals in the place of maxmin, and hence different ambiguity attitudes, may be obtained by modifying the Caution axiom in Corollary 1, in the style of Cerreia–Vioglio et al. (2025, Proposition 7); Lanzani (2025).

Under the individuals'-confidence interpretation (Section 3.1), the individuals' input can come in the form of confidence rankings, on which confidence aggregation operates directly. Just as pooling rules tacitly assume comparability of probability assignments across individuals, confidence aggregation requires interpersonal comparison of confidence. Hill (2019a) discusses the problem of 'calibrating' confidence levels across individuals, providing a scale for interpersonal confidence comparison which can be used for 'direct' applications of confidence aggregation. This comparison can alternatively be provided by preferences, in the context of the appropriate decision models (Sections 4.1 and 5.1).

Indeed, the characterisations in Sections 4 and 5 connect to a literature on belief aggregation working in preference-based frameworks. Spurious unanimity, for instance, first arose as an issue for preference aggregation with potentially differing utilities and subjective probabilities (Mongin, 1995, 2016), and only recently has been recognised as relevant for aggregation of belief *tout court*. For instance, Mongin and Pivato (2020); Dietrich (2021); Pivato (2022) criticise Gilboa et al.'s (2004) characterisation of utilitarian aggregation of utility and linear pooling of probabilities on these grounds. Several reactions in this literature work with preferences and restrict the domain of the Pareto condition. Dietrich (2021) restricts it to cases where all agents have identical subjective probabilities. Mongin and Pivato (2020); Pivato (2022) restrict Pareto to such an extent that their representations involve "no connection between the social probability and the individual ones". Unlike the aggregation approach developed here, these make no attempt to retain the consensus-preservation intuition behind Pareto. Bommier et al. (2026) present a condition preserving consensuses on prospects yielding identical distributions of outcomes for all individuals, and use it to provide a decision rule involving prospect-dependent aggregate distributions. Alon and Gayer (2016) and Stanca (2021) consider aggregation of EU preferences when group preferences may be non-expected utility, and under identical utilities in the latter case. Both involve versions of Pareto that, were group preferences expected utility, would lead to linear pooling.

To the extent that confidence rankings support both ambiguity averse and incomplete preferences (Sections 4.1 and 5.1), confidence aggregation provides an aggregation rule for both sorts of preferences. Crès et al. (2011) characterises an aggregation rule for maxmin-EU preferences, and Danan et al. (2016) explore aggregation of incomplete preferences with potentially differing utilities and beliefs. Both adopt conditions comparable with standard Pareto, whereas the approach proposed here leverages the non-probabilistic structure of beliefs in aggregation, building on the insight that confidence has a role in consensus formation (Section 2.2). That said, a possible direction for future research is to investigate whether the techniques used by Danan et al. (2016) to deal with differing utilities can extend our results to cases where the individuals and the decision maker have different tastes. Nau (2002) proposes an aggregation rule for a confidence-based belief representation which is a special case of that used here (Hill, 2016, Sect. 6). His rule is based on a different intuition, pertaining to the Bayesian risk function of the group, as defined in terms of an opponent's minimum expected loss in a betting game. Neither approach is contained in the other.[27]

[27] This can be seen from the fact that Nau's rule violates Eqn. (2); see Nau (2002,

# References

Abbas, A. E. (2009). A Kullback-Leibler view of linear and log-linear pools. *Decision Analysis*, 6(1):25–37.

Alon, S. and Gayer, G. (2016). Utilitarian Preferences With Multiple Priors. *Econometrica*, 84(3):1181–1201.

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

Bommier, A., Fabre, A., Goussebaile, A., and Heyen, D. (2026). Cautiousness when Experts Disagree. *Journal of Economic Theory*, page 106181.

Bordley, R. F. and Wolff, R. W. (1981). Note—on the aggregation of individual probability estimates. *Management Science*, 27(8):959–964.

Bradley, R. (2017a). *Decision theory with a human face*. Cambridge University Press.

Bradley, R. (2017b). Learning from others: conditioning versus averaging. *Theory and Decision*, pages 1–16.

Cerreia–Vioglio, S., Hansen, L. P., Maccheroni, F., and Marinacci, M. (2025). Making decisions under model misspecification. *Review of Economic Studies*, page rdaf046.

Chateauneuf, A. (1991). On the use of capacities in modeling uncertainty aversion and risk aversion. *Journal of Mathematical Economics*, 20(4):343–369.

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

Cooke, R. M. (1991). *Experts in uncertainty: opinion and subjective probability in science*. Oxford University Press.

Crès, H., Gilboa, I., and Vieille, N. (2011). Aggregation of multiple prior opinions. *Journal of Economic Theory*, 146:2563–2582.

Csiszar, I. (1975). *i*-divergence geometry of probability distributions and minimization problems. *The Annals of Probability*, 3(1):146–158.

Danan, E., Gajdos, T., Hill, B., and Tallon, J.-M. (2016). Robust Social Decisions. *The American Economic Review*, 106(9):2407–2425.

Dietrich, F. (2021). Fully Bayesian Aggregation. *Journal of Economic Theory*, 194:105255.

Fan, K. (1956). On systems of linear inequalities. Pages: 99–156 Publication Title: Linear inequalities and related systems Volume: 38.

French, S. (1985). Group consensus probability distributions: A critical survey in Bayesian statistics. *Bayesian statistics*, 2.

Genest, C. and Zidek, J. V. (1986). Combining Probability Distributions:

Figs 2 & 3).

A Critique and an Annotated Bibliography. *Statistical Science*, 1(1):114–135.

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

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

Gilboa, I., Samet, D., and Schmeidler, D. (2004). Utilitarian Aggregation of Beliefs and Tastes. *Journal of Political Economy*, 112(4):932–938.

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

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

Hansen, L. P. and Sargent, T. J. (2008). *Robustness.* Princeton university press.

Hansen, L. P. and Sargent, T. J. (2022). Structured ambiguity and model misspecification. *Journal of Economic Theory*, 199:105165.

Hill, B. (2013). Confidence and decision. *Games and Economic Behavior*, 82:675–692.

Hill, B. (2016). Incomplete preferences and confidence. *Journal of Mathematical Economics*, 65:83–103.

Hill, B. (2019a). Confidence in Belief, Weight of Evidence and Uncertainty Reporting. In *Proceedings of the Eleventh International Symposium on Imprecise Probabilities: Theories and Applications*, pages 235–245. PMLR.

Hill, B. (2019b). Confidence in Beliefs and Rational Decision Making. *Economics & Philosophy*, 35(2):223–258.

Jeffrey, R. C. (1972). *The Logic of Decision.* University of Chicago Press, Chicago, 2nd edition.

Karni, E. (2011). Continuity, completeness and the definition of weak preferences. *Mathematical Social Sciences*, 62(2):123–125.

Kemeny, J. G. (1959). Mathematics without numbers. *Daedalus*, 88(4):577–591.

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

Lanzani, G. (2025). Dynamic Concern for Misspecification. *Econometrica*, 93(4):1333–1370.

Liese, F. and Vajda, I. (1987). Convex statistical distances. *(No Title)*.

Luenberger, D. G. (1969). Optimization by vector space methods. Publisher: New York: John Wiley & Sons,.

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

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6):1022–1100.

Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S. L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M. I., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J. B. R., Maycock, T. K., Waterfield, T., Yelekçi, , Yu, R., and Zhou, B., editors (2021). *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

McConway, K. J. (1981). Marginalization and Linear Opinion Pools. *Journal of the American Statistical Association*, 76(374):410–414.

Mongin, P. (1995). Consistent Bayesian Aggregation. *Journal of Economic Theory*, 66(2):313–351.

Mongin, P. (2016). Spurious unanimity and the Pareto principle. *Economics & Philosophy*, 32(3):511–532.

Mongin, P. and Pivato, M. (2020). Social preference under twofold uncertainty. *Economic Theory*, 70(3):633–663.

Nau, R. F. (2002). The aggregation of imprecise probabilities. *Journal of Statistical Planning and Inference*, 105(1):265–282.

Pivato, M. (2022). Bayesian social aggregation with accumulating evidence. *Journal of Economic Theory*, 200:105399.

Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997). Bayesian model averaging for linear regression models. *Journal of the American Statistical Association*, 92(437):179–191.

Rockafellar, R. T. (1970). Convex Analysis (Princeton Mathematical Series). *Princeton University Press*, 46:49.

Stanca, L. (2021). Smooth aggregation of Bayesian experts. *Journal of Economic Theory*, 196:105308.

Steel, M. F. (2020). Model averaging and its use in economics. *Journal of Economic Literature*, 58(3):644–719.

Stone, M. (1961). The linear opinion pool. *Annals of Mathematical Statistics*, 32:1339–1342.

# Appendices: For Online Publication

## A Expertise-sensitive pooling: Running Example

In this Appendix, for completeness, we set out the details concerning the running Example in Sections 3.3–3.5, as well as some further comments and extensions to confidence in judgements about probabilistic independence.

### A.1 Examples 3.2–3.4

First note that each probability measure $p$ over $\{L, L^c\}$ is determined by $p(L)$, and similarly for the other partitions (Example 3.2). So each tuple $(p_L, p_R, p_B) \in \Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ is fully characterised by the vector $(p_L(L), p_R(R), p_B(B)) \in [0,1]^3$. The set $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B} \subseteq \Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ of tuples derived from probability measures on the full state space (Section 3.5) is defined by the following linear inequalities imposed by the fact that $B = L \cap R$: $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$ is the set of all $(p_L, p_R, p_B) \in \Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ satisfying

$$\begin{aligned} p_L(L) &\geqslant p_B(B) \\ p_R(R) &\geqslant p_B(B) \\ 1 &\geqslant p_L(L) + p_R(R) - p_B(B) \end{aligned} \tag{A.1}$$

Writing these inequalities in vector notation, $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$ corresponds to the set of vectors $\mathbf{q} \in [0,1]^3$ satisfying the constraint $\mathbf{A}\mathbf{q} \leqslant \mathbf{r}$ where

$$\mathbf{A} = \begin{pmatrix} -1 & 0 & 1 \\ 0 & -1 & 1 \\ 1 & 1 & -1 \end{pmatrix}, \quad \mathbf{r} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Moreover, for each confidence level $o$ in the confidence ranking (5) in Example 3.3, the map of $c^L(o)$ into the space $\Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ can be written as:

$$c^L(o) = \{\mathbf{q} \in [0,1]^3 : (\mathbf{q} - \mathbf{p^L})^T \mathbf{D^L} (\mathbf{q} - \mathbf{p^L}) \leqslant o\} \tag{A.2}$$

where

$$\mathbf{p^L} = \begin{pmatrix} 0.9 \\ 0.1 \\ 0.09 \end{pmatrix}, \quad \mathbf{D^L} = \begin{pmatrix} w_L^L & 0 & 0 \\ 0 & w_R^L & 0 \\ 0 & 0 & w_B^L \end{pmatrix}$$

and similarly for Ray, with

$$\mathbf{p^R} = \begin{pmatrix} 0.1 \\ 0.9 \\ 0.09 \end{pmatrix}, \quad \mathbf{D^R} = \begin{pmatrix} w_L^R & 0 & 0 \\ 0 & w_R^R & 0 \\ 0 & 0 & w_B^R \end{pmatrix}$$

It follows that the minimisation problem (8) defining the centre of the confidence aggregation in Example 3.4 becomes the following simple quadratic optimisation problem under constraints:

$$\underset{\mathbf{Aq} \leqslant \mathbf{r}}{\arg\min} \sum_{i=L,R} (\mathbf{q} - \mathbf{p^i})^T \mathbf{D^i} (\mathbf{q} - \mathbf{p^i})$$

If $\left(\frac{w_L^L}{w_L^L + w_L^R} p^L(L) + \frac{w_L^R}{w_L^L + w_L^R} p^R(L)\right) + \left(\frac{w_R^L}{w_R^L + w_R^R} p^L(R) + \frac{w_R^R}{w_R^L + w_R^R} p^R(R)\right) - 1 \leqslant 0.09$, then the constraints are slack, and the unqiue solution is:

$$\begin{aligned} p(L) &= \frac{w_L^L}{w_L^L + w_L^R} p^L(L) + \frac{w_L^R}{w_L^L + w_L^R} p^R(L) \\ p(R) &= \frac{w_R^L}{w_R^L + w_R^R} p^L(R) + \frac{w_R^R}{w_R^L + w_R^R} p^R(R) \\ p(B) &= 0.09 \end{aligned}$$

Henceforth, denote the issue-wide weighted averages by $m_L = \frac{w_L^L p^L(L) + w_L^R p^R(L)}{w_L^L + w_L^R}$

and $m_R = \frac{w^L_R p^L(R) + w^R_R p^R(R)}{w^L_R + w^R_R}$. Otherwise, solving the minimisation problem yields the following unique solution:

$$\begin{aligned}
p(L) =& m_L - \frac{1}{w^L_L + w^R_L} \frac{m_L + m_R - 1.09}{\frac{1}{w^L_B + w^R_B} + \frac{1}{w^L_R + w^R_R} + \frac{1}{w^L_L + w^R_L}} \\
p(R) =& m_R - \frac{1}{w^L_R + w^R_R} \frac{m_L + m_R - 1.09}{\frac{1}{w^L_B + w^R_B} + \frac{1}{w^L_R + w^R_R} + \frac{1}{w^L_L + w^R_L}} \\
p(B) =& 0.09 + \frac{1}{w^L_B + w^R_B} \frac{m_L + m_R - 1.09}{\frac{1}{w^L_B + w^R_B} + \frac{1}{w^L_R + w^R_R} + \frac{1}{w^L_L + w^R_L}}
\end{aligned}$$

where, as specified, $p^L(L) = p^R(R) = 0.9$, $p^L(R) = p^R(L) = 0.1$, $p^L(B) = p^R(B) = 0.09$. These solutions yield Proposition 4, as follows.

*Proof of Proposition 4.* The solutions provided above imply that the centre of the aggregate confidence ranking is a singleton. We now consider the other parts of the Proposition.

*Part i.* If the constraint mentioned above does not bind (i.e. $\left(\frac{w^L_L}{w^L_L + w^R_L} p^L(L) + \frac{w^R_L}{w^L_L + w^R_L} p^R(L)\right) + \left(\frac{w^L_R}{w^L_R + w^R_R} p^L(R) + \frac{w^R_R}{w^L_R + w^R_R} p^R(R)\right) - 1 \leqslant 0.09$), then the result follows immediately from the assumptions (notably $w^L_L > w^R_L$ and $w^R_R > w^L_R$) and the solution provided above. Now consider the case where the constraint is binding. Since $\frac{w^R_L}{w^L_L} = \frac{w^L_R}{w^R_R}$ (and plugging in the values for $p^L(L), p^R(L), p^L(R), p^R(R)$), we have that $m_L = m_R$. Using the solutions provided above, and the fact that the assumption implies that $(w^L_L + w^R_L)\left(\frac{1}{w^L_B + w^R_B} + \frac{1}{w^L_R + w^R_R} + \frac{1}{w^L_L + w^R_L}\right) \geqslant 2$, so

$$m_L - \frac{1}{w^L_L + w^R_L} \frac{m_L + m_R - 1.09}{\frac{1}{w^L_B + w^R_B} + \frac{1}{w^L_R + w^R_R} + \frac{1}{w^L_L + w^R_L}} \geqslant m_L - (m_L - 0.545) = 0.545$$

and similarly for the expression for $p(R)$, yielding the required result.

*Part ii.* The result follows from taking the limits in the previous solutions as $\frac{w^L_B + w^R_B}{w^L_L + w^R_L} \to 0$ and $\frac{w^L_B + w^R_B}{w^L_R + w^R_R} \to 0$.

□

Note that Proposition 4 considers the case where $\frac{w^L_B + w^R_B}{w^L_L + w^R_L} \to 0$ and $\frac{w^L_B + w^R_B}{w^L_R + w^R_R} \to 0$, which, arguably, is closest to the example described in the Introduction. For completeness, note that, in the opposite case of $\frac{w^L_B + w^R_B}{w^L_L + w^R_L} \to$

$\infty$ and $\frac{w_B^L+w_B^R}{w_R^L+w_R^R} \to \infty$, the confidence in the probability judgements concerning $B$ grows very large comparatively, so these are retained at the expense of others. Hence, we have:

$$\begin{aligned} p(L) \to & m_L - \frac{m_L + m_R - 1.09}{1 + \frac{w_L^L + w_L^R}{w_R^L + w_R^R}} \\ p(R) \to & m_R - \frac{m_L + m_R - 1.09}{1 + \frac{w_R^L + w_R^R}{w_L^L + w_L^R}} \\ p(B) \to & 0.09 \end{aligned}$$

Here the judgement about $B$ is fully preserved, as one would expect given the high confidence postulated in it. This places a strong constraint on $p(L)$ and $p(R)$ (namely, $p(L) + p(R) = 1.09$). The possible probability available is shared between $L$ and $R$ according to the comparison between the issue-wide weighted averages $m_L$ and $m_R$ and the ratio between the overall confidence (i.e. $w_L^L + w_L^R$ v.s. $w_R^L + w_R^R$) in each of these judgements.

## A.2 Confidence in independence judgements

A central factor in Example 3.4 is the trade-off between the confidence in the judgements concerning the main two issues—labour and real estate—and what happens to both, considered as a third issue. However, an alternative analysis considers individuals to have opinions on the main issues and their relationship, rather than 'primitive' views on $B$. We now briefly show that the confidence approach can easily support such perspectives.

**Example A.1.** Now suppose that Laura and Ray hold beliefs about $L$ and $R$, and about the independence of $L$ and $R$: they believe them to be independent, without being maximally confident in this judgement.[28] This can be reflected using a vector of weights $\mathbf{w^L} = (w_L^L, w_R^L, w_I^L)$ (resp.

[28]Independence here refers to the probabilistic sense: $p^i(L \cap R) = p^i(L)p^i(R)$. Note that the belief in independence implies that $p^L(B) = p^R(B) = 0.09$, as per Table 1.

$\mathbf{w^R} = (w^R_L, w^R_R, w^R_I))$ and the following confidence ranking:

$$c^L_{Ind}(o) = \left\{ q \in \Delta : \begin{array}{c} \sum_{j=\{L,R\}} w^L_j (q(j) - p^L(j))^2 \\ +w^L_I \left(q(B) - q(L).q(R)\right)^2 \end{array} \leqslant o \right\} \tag{A.3}$$

and similarly for $c^R_{Ind}$. These are clearly well-defined confidence rankings. The weighted element corresponding to the event $B$ here is $(q(B) - q(L).q(R))^2$, which reflects the 'distance' from independence of $L$ and $R$. So, at higher confidence levels, probability measures with larger 'distances' from independence are contained in the set of priors, translating the limited confidence in independence.

The solution of the minimisation problem (8) can be obtained similarly to the analysis in Example 3.4, yielding as centre of the aggregate confidence ranking $p$ with:

$$\begin{aligned} p(L) =& \frac{w^L_L}{w^L_L + w^R_L} p^L(L) + \frac{w^R_L}{w^L_L + w^R_L} p^R(L) \\ p(R) =& \frac{w^L_R}{w^L_R + w^R_R} p^L(R) + \frac{w^R_R}{w^L_R + w^R_R} p^R(R) \\ p(B) =& p(L).p(R) \end{aligned}$$

Here the aggregation on each of the issues $L$ and $R$ uses issue-specific weights, reflecting differing confidence, as in the limit case in Proposition 4. For the issue $B$, agents' beliefs concerning the independence of $L$ and $R$ generates the probability.

In tandem with the discussion in Section 3, this illustrates that the confidence approach can not only recoup averaging with issue-specific weights whilst retaining consistency, but it can also incorporate varying opinions about independence or more generally the relationship between issues. This is relevant for another recurrent criticism of linear pooling: that it does not preserve independence. As is well known, even if all individuals consider the events $L$ and $R$ to be independent, the linear pool might not (e.g. Genest and Zidek, 1986). This is easy to see on our running example: the linear pool of Laura's and Ray's probabilities with equal weights ($w^L = \frac{1}{2}$) is $p^{LP}(L) = 0.5$, $p^{LP}(R) = 0.5$, $p^{LP}(B) = 0.09$, so $L$ and $R$ are not independent under $p^{LP}$, though they are under $p^L$ and $p^R$. The aggregation above based on confidence rankings of the form (A.3) shows how confidence aggregation can respect independence, whilst retaining much of the spirit of linear pooling. For instance, when $w^L_L = w^L_R = w^R_L = w^R_R$, the resulting centre probability is $p(L) = 0.5$, $p(R) = 0.5$, $p(B) = 0.25$: i.e. the same as linear pooling for the issues $L$ and $R$, but with independence retained (and hence a different $B$).

The beliefs about the independence of $L$ and $R$ in Example A.1 are considered merely for the purposes of illustration. The point of the example is more general: by incorporating conditional probabilities much in the way proposed in Eq. (A.3), the confidence approach can respect conditional probability judgements (including, but not limited to, judgements about independence) in the aggregate belief. In accordance with the philosophy behind the approach, they are respected to the extent that the individuals are confident in them.

# B Characterising confidence aggregation: special cases

In this Appendix, we extend Theorem 1 to characterise, as special cases, confidence aggregation under generalisations of some of the families of confidence-level aggregators mentioned in Section 2.2 (Table 3).

**Example B.1** (Affine aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \max\{\sum_{i=1}^{n} \lambda_i o_i + \chi, 0\}$ for $\lambda_i \in \mathbb{R}_{\geqslant 0}$, $\chi \in \mathbb{R}$.

**Example B.2** (Generalised Average aggregator)**.** An affine aggregator for which $\lambda_i = \frac{1}{n}$ for all $i$.

**Example B.3** (Generalised Maximum aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \psi(\max\{o_i\})$ for $\psi : O \to O$ an increasing transformation of confidence levels.

**Example B.4** (Generalised Minimum aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \psi(\min\{o_i\})$ for $\psi : O \to O$ an increasing transformation of confidence levels.

**Example B.5** (Non-neutral Maximum aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \max\{\psi_i(o_i)\}$, where $\psi_i : O \to O$ (for $i = 1, \ldots, n$) are increasing transformations of confidence levels.

More specifically, we will provide results for the following stronger representation: $c^{rep} = \dot{F}_{\otimes}(c^1, \ldots, c^n)$, with $\dot{F}_{\otimes}(c'^1, \ldots, c'^n) = \overline{\dot{\Phi}(c'^1, \ldots, c'^n)}$ for every $(c'^1, \ldots, c'^n) \in \Pi^n$, where, for every $o \in O$ such that $\bigcup_{\mathbf{o}:\otimes\mathbf{o}=o} \bigcap_i c'^i(o_i) \neq \varnothing$

$$\dot{\Phi}_{\otimes}(c^1, \ldots, c^n)(o) = \bigcup_{\mathbf{o}:\otimes\mathbf{o}=o} \bigcap_i c'^i(o_i) \tag{B.1}$$

The only difference with respect to Definition 3 is that here the union is taken over all tuples of confidence levels whose confidence-level aggregate equals $o$, whereas the previous procedure looks at all those with confidence-level aggregate at most $o$. It follows directly from the fact that confidence rankings are increasing in $o$ that, if $c^{rep} = \dot{F}_{\otimes}(c^1, \ldots, c^n)$, then $c^{rep}$ is a consensus-preserving confidence aggregation in the sense of Definition 3.

Recall that, under Assumption 1, $\max\{D(\bullet), D(\bullet)\}$ is a monotonically increasing transformation of $\sigma$. By appropriate choice of normalisation for $O$ and $\mathcal{S}$, it can be assumed that they are identical. Under this assumption, we have the following result, which involves the axioms in Figure 4, and defines clauses according to Table 6.

**Theorem B.1.** *Suppose that $O$ is infinite, let $\{\succ^J\}_{J\subseteq I}$ satisfy Assumption 1 with $\max\{D(f), D(g)\} = \sigma(f, g)$ for all $f, g \in \mathcal{A}$. For each of the rows in Table 6: $\{\succ^J\}_{J\subseteq I}$ satisfy Consensus-wise Pareto, Consensus-based preferences, Non-degeneracy and the axiom(s) in the first column of the table if and only if, for each $J$ with $|J| \geqslant 2$, there exists a confidence-level aggregator $\otimes$ of the type specified in the second column such that $c^J = \dot{F}_{\otimes}(c^1, \ldots, c^n)$, up to convex closure.*

We make no particular claim for any of the confidence-level aggregators in Table 6 on normative grounds; we present this result to illustrate the richness of the approach, and exemplify some simple aggregators.

The axiom involved in the characterisation of confidence aggregation with an affine aggregator, Consensus Independence, uses the notion of uncovered consensus. For every tuple of stakes levels $\mathbf{s}$ exhibiting consensus

| Axioms | Aggregator |
|---|---|
| Consensus Independence | Affine |
| Consensus Independence, Neutrality | Generalised Average |
| Consensus Join | Non-neutral Maximum |
| Consensus Join, Neutrality | Generalised Maximum |
| Consensus Meet, Neutrality | Generalised Minimum |

Table 6: Characterisations of special case confidence-level aggregators, to be read in the context of Theorem B.1.

**Axiom** (Consensus Independence). *For all non-empty $J \subseteq I$, and all tuples of stakes levels $\mathbf{s_1}, \ldots, \mathbf{s_l}, \mathbf{t_1}, \ldots, \mathbf{t_m} \in \mathcal{S}^{|J|}$ exhibiting consensus and $\alpha_1, \ldots, \alpha_l, \beta_1, \ldots, \beta_m \in [0,1]$ with $\sum_{k=1}^{l} \alpha_k = \sum_{j=1}^{m} \beta_j = 1$ and $\sum_{k=1}^{l} \alpha_k \mathbf{s}_k = \sum_{j=1}^{m} \beta_j \mathbf{t}_j$, if, for some stakes levels $s_1, \ldots, s_l, t_1, \ldots, t_m$, $\succ^J$ does not respect the consensuses $\succ_{\mathbf{s}_k}$ at $s_k$ for each $k = 1, \ldots, l$, and $\succ_{\mathbf{t}_j}$ are uncovered consensuses at $t_j$ for all $j = 1, \ldots, m$, then $\sum_{k=1}^{l} \alpha_k s_k < \sum_{j=1}^{m} \beta_j t_j$.*

**Axiom** (Consensus Join). *For all non-empty $J \subseteq I$ and any set of tuples of stakes levels $\{\mathbf{s}_k : k \in K\}$ exhibiting consensus, if $\succ^J$ respects the consensuses $\succ_{\mathbf{s}_k}$ at $s$ for all $k \in K$, then it respects the consensus $\succ_{\bigvee_{k \in K} \mathbf{s}_k}$ at $s$.*

**Axiom** (Consensus Meet). *For all non-empty $J \subseteq I$ and any set of tuples of stakes levels $\{\mathbf{s}_k : k \in K\}$ exhibiting consensus, if $\succ^J$ respects the consensuses $\succ_{\mathbf{s}_k}$ at $s$ for al l $k \in K$, then it respects the consensus $\succ_{\bigwedge_{k \in K} \mathbf{s}_k}$ at $s$.*

**Axiom** (Neutrality). *For all non-empty $J \subseteq I$ and any stakes levels $s$, tuple of stakes levels $\mathbf{s}$ and permutation $\pi$ such that $\mathbf{s}, \pi(\mathbf{s})$ exhibit consensus, $\succ^J$ respects the consensus $\succ_{\mathbf{s}}$ at $s$ if and only if $\succ^J$ respects the consensus $\succ_{\pi(\mathbf{s})}$ at $s$.*

Where, for any $\mathbf{s}, \mathbf{t}, \mathbf{s}_k \in \mathcal{S}^n$ for $k \in K$ and $\alpha \in [0,1]$, $(\alpha \mathbf{s} + (1-\alpha)\mathbf{t})_i = \alpha s_i + (1-\alpha) t_i$, $(\bigvee_{k \in K} \mathbf{s})_i = \sup\{(s_k)_i : k \in K\}$ and $(\bigwedge_{k \in K} \mathbf{s})_i = \inf\{(s_k)_i : k \in K\}$.

Figure 4: Axioms for special cases

and stakes level $s$ with $\succ^J$ respecting the consensus $\succ_{\mathbf{s}}$ at $s$, we say that the consensus at $s$ is *covered* when, for all acts $f, g$, if $f \nsucc_{\mathbf{s}} g$ then there exists a tuple $\mathbf{s}'$ exhibiting consensus with $\mathbf{s}' \ngeq \mathbf{s}$ such that $\succ^J$ respects the consensus $\succ_{\mathbf{s}'}$ at $s$ and $f \nsucc_{\mathbf{s}'} g$. Otherwise, say that the consensus is *uncovered* at $s$. When the consensus $\succ_{\mathbf{s}}$ is covered, there is no $f, g$ such that an absence of preference between them according to $\succ^J_s$ can be pinpointed as being due to the respect for consensus $\succ_{\mathbf{s}}$, for there is some other consensus respected at $s$ that does not have any preference either. So, when the consensus is uncovered, it contributes for sure to the construction of group preferences, even in the context of the other relevant consensuses. In particular, it means that the group confidence level assigned to this consensus can't be a lower than that corresponding to stakes level $s$.

In the light of this, Consensus Independence can be thought of as an Independence-like axiom, adapted to this context. An Independence axiom in this context would imply that if $\succ^J$ does not respect $\succ_{\mathbf{s}_i}$ at $s_i$, for all $i$, then it does not respect any mixture $\succ_{\sum_k \alpha_k \mathbf{s}_k}$ exhibiting consensus at $\sum_k \alpha_k s_k$. However, consensus-preserving aggregation with an affine aggregator can violate such a condition when, for instance, the consensus involved is respected 'by accident', because it is covered; so the implication does not hold. Consensus Independence corrects the first-pass independence condition to account for such cases, using the notion of uncovered consensus. It allows that the mixture of uncovered consensuses may not be uncovered, and it allows that a mixture of non-respected consensus may be respected, but doesn't allow the mixture of uncovered consensuses to coincide with a mixture of non-respected ones.

The characterising axiom for a non-neutral maximum confidence-level aggregator, Consensus Join, states that respect for consensus at $s$ is preserved if one takes the consensus corresponding to the largest stakes level for each entry in the tuple (the join). Consensus Meet is the dual axiom, involving the lowest stakes level for each entry. Neutrality is a standard neutrality axiom, adapted to the current context, stating that respect for consensus is preserved under permutation of individuals. Added to the other conditions, it characterises the 'neutral' average, generalised maximum and generalised minimum confidence-level aggregators.

# C Connecting Confidence-based and Variational Preferences

We briefly state two results mentioned in Section 5.1, connecting the ambiguity averse confidence-based preferences (11) with variational preferences.

A preference $\succeq$ is *variational* (Maccheroni et al., 2006) if it is represented by

$$V(f) = \min_{p \in \Delta^{fin}} \left(\mathbb{E}_p u(f) + \gamma(p)\right) \tag{C.1}$$

where $u$ is as in Section 4.1 and $\gamma = \Delta^{fin} \to [0, \infty]$ is a grounded, convex, lower semicontinuous function on the space on the space of finitely additive probability measures.[29] The preference is *unbounded continuous variational* (Maccheroni et al., 2006) if $u$ is unbounded (above or below) and the Monotone Continuity axiom holds, in which case (C.1) holds with the space of countably additive probability measures instead of $\Delta^{fin}$ (Maccheroni et al., 2006, Sect 3.4). As noted by Maccheroni et al. (2006), the countably additive case is relevant for variational preferences generated by divergences, including from sets of models (Cerreia–Vioglio et al., 2025).

The following result shows that both classes of variational preferences are in fact confidence preferences, for appropriate $c$ and $D$ (see Appendix D.3 for proofs).

**Proposition C.1.** *Let $\succeq$ be a variational preference or an unbounded continuous variational preference. Then there exists $(c, D, u)$ such that $\succeq$ is represented by $(c, D, u)$ according to* (11).

For a partial converse to this result, let $\succ^{dm}$ be a preference relation represented by $(c, D, u)$ according to (11), and consider the following two assumptions on $D$.

**Assumption C.1.** *For all $f \in \mathcal{A} \backslash \mathcal{A}^c$, $x, y \in \mathcal{A}^c$ and $\alpha \in (0, 1]$*

$$D(f_\alpha x) = D(f_\alpha y) \tag{C.2}$$

[29] $\gamma$ is grounded if its infimum is zero: as elsewhere, $\Delta$ is taken to be equipped with the weak* topology.

**Assumption C.2.** *For all $f, g \in \mathcal{A}$, $\alpha \in [0,1]$, $x \in \mathcal{A}^c$ and $o \in O$, if $f \succ^{dm} x$, $g \succ^{dm} x$ and $\min_{p \in c(o)} \mathbb{E}_p u(f_\alpha g) \leqslant u(x)$, then*

$$D(f_\alpha g) < o$$

Finally, a confidence ranking $c$ is said to be strongly increasing if, for every $o, o'$, $o < o'$, every boundary point of $c(o)$ is not a boundary point of $c(o')$.

We have the following result.

**Proposition C.2.** *Let $\succeq^{dm}$ be represented by $(c, D, u)$ according to (11). If $D$ satisfies Assumptions C.1 and C.2, then $\succeq^{dm}$ is a variational preference. Moreover, if $c$ is strongly increasing, the converse implication holds.*

So the confidence preferences used in (11) can coincide with variational preferences, under appropriate assumptions on $D$. Under the philosophy behind confidence preferences (Hill, 2019b), $D$ typically reflects *only* tastes about the confidence level deemed appropriate for evaluating a given act, and hence the stakes involved in it. Such an interpretation is difficult to uphold for Assumption C.2, suggesting that the subfamily of confidence preferences that are variational do not respect this aspect of the taste-belief separation. This may be related to the lack of such separation for variational preferences in general, as suggested by the dependence of $\gamma$ on $u$ in the model's uniqueness result (Maccheroni et al., 2006).

# D Proofs

## D.1 Proofs of results in Sections 2, 4 and Appendix B

We begin with the following Proposition, mentioned in Section 2.2.

**Proposition D.1.** *$c^{rep}$ is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$ under $\otimes$ if and only if*

$$\iota^{rep}(p) = \begin{cases} 0 & p \in \min_{p' \in \Delta} \otimes(\iota^1(p'), \ldots, \iota^n(p')) \\ \otimes(\iota^1(p), \ldots, \iota^n(p)) & \textit{otherwise.} \end{cases}$$

*where the implausibility function $\iota^i$ for $c^i$ is as defined in Section 2.1.*

*Proof.* By the definition, $p \in c^{rep}(o)$ if and only if, for some $\mathbf{o}$ with $\otimes\mathbf{o} \leqslant o$, $p \in c^i(o_i)$ for all $i = 1, \ldots, n$. For any $p \in \Delta$, if $\iota^i(p) \in O$, then $p \in c^i(\iota^i(p))$, whereas if $\iota^i(p) = \infty$, then $p \notin c^i(o)$ for all $o \in O$. If $\otimes(\iota^1(p), \ldots, \iota^n(p)) = \infty$ the result follows immediately from the definitions; henceforth, consider $p$ such that $\otimes(\iota^1(p), \ldots, \iota^n(p)) \in O$. For such $p$, $p \in c^i(\iota^i(p))$ for all $i$ such that $\iota^i(p) \in O$, and hence $p \in c^{rep}(\otimes(\iota^1(p), \ldots, \iota^n(p)))$. If there exists no $p' \in \Delta$ with $\otimes(\iota^1(p'), \ldots, \iota^n(p')) < \otimes(\iota^1(p), \ldots, \iota^n(p))$, then $p \in \bigcap_{o:\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i c'^i(o_i)\neq\varnothing} \bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i c'^i(o_i)$, so $p \in c^{rep}(0)$ by Definition 3. Hence $\iota^{rep}(p) = 0$, yielding the first clause of the required formula. Now consider the case where there exists $p' \in \Delta$ with $\otimes(\iota^1(p'), \ldots, \iota^n(p')) < \otimes(\iota^1(p), \ldots, \iota^n(p))$. For any $\mathbf{o}$ with $\otimes\mathbf{o} < \otimes(\iota^1(p), \ldots, \iota^n(p))$, $o_i < \iota^i(p)$ for some $i$ with $\iota^i(p) \in O$ by the monotonicity of $\otimes$; since $\iota^i(p) = \min\{o' \in O : p \in c^i(o')\}$, it follows that $p \notin c^i(o_i)$. Hence, for every $o' < \otimes(\iota^1(p), \ldots, \iota^n(p))$ such that $\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o'}\bigcap_i c'^i(o_i) \neq \varnothing$, $p \notin c^{rep}(o')$. Since in the current case, there exists such $o' = \otimes(\iota^1(p'), \ldots, \iota^n(p'))$, the second clause of the required formula follows from the definition of $\iota$. $\square$

We now prove Theorems 1 and B.1. It suffices to prove them for a single $J \subseteq I$, $|J| \geqslant 2$. For the rest of this section, we fix such a $J$, and use $\succ^0$ to denote $\succ^J$, and $c^0$ to denote $c^J$. Recall that, under Assumption 1, $\{(c^i, D, u)\}, (c^0, D, u)$ denote the representations of the $\{\succ^i\}, \succ^0$. Moreover, as noted in Section 4.1 (footnote 19), $\max\{D(f), D(g)\} = \zeta \circ \sigma$ for some strictly increasing $\zeta : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$. Throughout the rest of this section, with slight abuse of notation, for any stakes level $s \in \mathcal{S}$, we shall denote $c^i(\zeta(s))$ by $c^i(s)$, for all $i$.

### D.1.1 Proof of Theorem 1

We first show sufficiency of the axioms. Let $X \subseteq \mathcal{S}^n$ be the set of tuples exhibiting consensus. By Non-degeneracy, there is a tuple $(s_i)_{i\in I}$ exhibiting consensus, so there is a tuple $(s_i)_{i\in J}$ exhibiting consensus for any $J \subseteq I$, $|J| \geqslant 2$. Hence $X \neq \varnothing$. Let $\geqslant$ be the dominance ordering on $\mathcal{S}^n$: $\mathbf{s} \geqslant \mathbf{t}$ if and only if $s_i \geqslant t_i$ for all $i$. $X$ is closed under $\geqslant$: if $\mathbf{s} \in X$ and $\mathbf{t} \geqslant \mathbf{s}$, then $\succ_{\mathbf{s}}$ is not contradictory; but $\succ_{t_i} \subseteq \succ_{s_i}$ for all $i$ by the properties of confidence rankings, so $\succ_{\mathbf{t}}$ is not contradictory and hence $\mathbf{t} \in X$.

The following claim follows immediately from standard arguments (e.g. Ghirardato et al., 2004), for every $\succ_{\mathbf{s}}$ exhibiting consensus.

**Claim D.1.** $\succ^0$ respects the consensus $\succ_{\mathbf{s}}$ at stakes level $s$ if and only if $c^0(s) \supseteq \bigcap_{i=1}^n c^i(s_i)$.

**Claim D.2.** For any set $Y \subseteq \mathcal{S}^n$ such that $\succ_{\mathbf{s}}$ exhibits consensus for every $\mathbf{s} \in Y$, $\bigcap_{\mathbf{s}\in Y} \succ_{\mathbf{s}}$ is represented by $\bigcup_{\mathbf{s}\in Y} \bigcap_{i=1}^n c^i(s_i)$ in the following sense: for all $f, g \in \mathcal{A}$, $f \succ_{\mathbf{s}} g$ if and only if

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \quad \text{for all } p \in \bigcup_{\mathbf{s}\in Y} \bigcap_{i=1}^n c^i(s_i) \tag{D.1}$$

*Proof of Claim D.2.* First consider $\mathbf{s}$ exhibiting consensus, and let $\succ_{\bigcap \mathbf{s}}$ be the Bewley preference such that, for every $f, g \in \mathcal{A}$, $f \succ_{\bigcap \mathbf{s}} g$ if and only if

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \quad \text{for all } p \in \bigcap_{i=1}^n c^i(s_i) \tag{D.2}$$

Note that, since the $c^i$ are closed and convex, so is their intersection. For every $f, g \in \mathcal{A}$, if $f \succ_{s_i}^i g$ for some $i$ and $f \not\prec_{s_i}^i g$ for every $i$, then, by Assumption 1, for some $i$, $\mathbb{E}_p u(f) > \mathbb{E}_p u(g)$ for all $p \in c^i(s_i)$, and, for every $i$, it is not the case that $\mathbb{E}_p u(f) < \mathbb{E}_p u(g)$ for all $p \in c^i(s_i)$. Since $\bigcap_{i=1}^n c^i(s_i) \neq \varnothing$, this holds if and only if, for all $p \in \bigcap_{i=1}^n c^i(s_i)$, $\mathbb{E}_p u(f) > \mathbb{E}_p u(g)$. Hence $\bigcup_{i=1}^n \succ_{s_i}^i \subseteq \succ_{\bigcap \mathbf{s}}$. Now take any Bewley preference $\succ'$ extending $\bigcup_{i=1}^n \succ_{s_i}^i$. It follows that, for each $i$, $\succ'$ extends $\succ_{s_i}^i$, so, for every $f, g \in \mathcal{A}$, if there exists some $i$ such that $\mathbb{E}_p u(f) > \mathbb{E}_p u(g)$ for all $p \in c^i(s_i)$, then $f \succ' g$. Since $\succ'$ is a Bewley preference represented by a set of priors $\mathcal{C}$, it follows that $\mathcal{C} \subseteq \bigcap_{i=1}^n c^i(s_i)$. Hence $\succ' \supseteq \succ_{\bigcap \mathbf{s}}$, whence, by the definition of $\succ_{\mathbf{s}}$, $\succ_{\mathbf{s}} = \succ_{\bigcap \mathbf{s}}$.

Now consider $Y$ as specified. The case in which $Y$ is a singleton has just been treated, so suppose that $Y$ contains several elements. By the previous observation, for every $f, g \in \mathcal{A}$, $f \succ_{\mathbf{s}} g$ for every $\mathbf{s} \in Y$ if and only if $f \succ_{\bigcap \mathbf{s}} g$ for every $\mathbf{s} \in Y$, which holds if and only if $\mathbb{E}_p u(f) > \mathbb{E}_p u(g)$ for all $p \in \bigcap_{i=1}^n c^i(s_i)$ for every $\mathbf{s} \in Y$. This holds if and only if $\mathbb{E}_p u(f) > \mathbb{E}_p u(g)$ for all $p \in \bigcup_{\mathbf{s}\in Y} \bigcap_{i=1}^n c^i(s_i)$, as required.

□

Define the function $G : X \to \mathcal{S}$ as follows:

$$\begin{aligned} G(\mathbf{s}) &= \min\left\{s : \succ^0_s \subseteq \succ_{\mathbf{s}}\right\} \\ &= \min\left\{s : c^0(s) \supseteq \bigcap_{i=1}^{n} c^i(s_i)\right\} \end{aligned}$$

where the equality follows from Claim D.1. Note that if $G(X)$ is a finite set, then $\min G(X) \in G(X)$. Whilst this might not be the case for infinite $O$, the following Proposition implies that $X$ is closed. As a point of terminology, a confidence ranking $c$ is said to be strictly increasing just below $s$ if, for all $s' < s$, $c(s') \subsetneq c(s)$.

**Proposition D.2.** *For any decreasing sequence $\mathbf{s}^j \in X$ with $\mathbf{s}^j \to \mathbf{s}$, $\mathbf{s} \in X$. Moreover, $G(\mathbf{s}) \leqslant \lim G(\mathbf{s}^j)$ with equality if $c^0$ is strictly increasing just below $\lim G(\mathbf{s}^j)$.*

*Proof.* Consider a decreasing sequence $\mathbf{s}^j \in X$ with $\mathbf{s}^j \to \mathbf{s}$. Since each $c^i$ is upper semicontinuous, $\bigcap_j c^i(s_i^j) = c^i(s_i)$ for each $i$, so $\bigcap_{i=1}^n c^i(s_i) = \bigcap_{i=1}^n \bigcap_j c^i(s_i^j) = \bigcap_j \bigcap_{i=1}^n c^i(s_i^j) \neq \varnothing$. So $\mathbf{s} \in X$. Moreover, since $\mathbf{s}^j$ is decreasing, $\bigcap_{i=1}^n c^i(s_i^j)$ is decreasing, and hence $G(\mathbf{s}^j)$ is decreasing. Consider $s = \lim G(\mathbf{s}^j)$; by the upper semicontinuity of $c^0$, $c^0(s) = \bigcap_j c^0(G(\mathbf{s}^j)) \supseteq \bigcap_j \bigcap_{i=1}^n c^i(s_i^j)$. Hence, by the definition of $G$, $G(\mathbf{s}) \leqslant s$. Moreover, for every $s' < s$ with $c^0(s') \subsetneq c^0(s)$, there exists $j$ such that $c^0(s') \nsupseteq \bigcap_{i=1}^n c^i(s_i^j)$, so $G(\mathbf{s}) > s'$. Hence $G(\mathbf{s}) \leqslant \lim G(\mathbf{s}^j)$ with equality if $c^0$ is strictly increasing just below $\lim G(\mathbf{s}^j)$, as required. □

**Claim D.3.** For every $s \in G(X)$, $\succ^0_s$ is represented by $\bigcup_{\mathbf{s} \in X : s \geqslant G(\mathbf{s})} \bigcap_i c^i(s_i)$ according to D.2: i.e. for all $f, g \in \mathcal{A}$, $f \succ^0_s g$ if and only if:

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \quad \text{for all } p \in \bigcup_{\mathbf{s} \in X : s \geqslant G(\mathbf{s})} \bigcap_i c^i(s_i) \tag{D.3}$$

*Proof.* Fix a stakes level $s \in G(X)$ and consider any $\mathbf{s}'$ with $G(\mathbf{s}') \leqslant s$. By the definition of $G$, there exists $s'' \in \mathcal{S}$ with $s'' \leqslant s$ and $\succ^0_{s''} \subseteq \succ_{\mathbf{s}'}$. It follows from the nestedness properties of confidence rankings that $\succ^0_s \subseteq \succ^0_{s''} \subseteq \succ_{\mathbf{s}'}$. Since this holds for all $\mathbf{s}'$ with $G(\mathbf{s}') \leqslant s$, it follows that $\succ^0_s \subseteq \bigcap_{\mathbf{s} \in X : s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$.

To establish the opposite containment, consider $f, g$ with $f \succ_{\mathbf{s}} g$ for all $\mathbf{s} \in X$ with $s \geqslant G(\mathbf{s})$. For any $\mathbf{s}'$ such that $\succ^0$ respects the consensus $\succ_{\mathbf{s}'}$

at $s$, it follows from the definition of $G$ that $s \geqslant G(\mathbf{s}')$, so $f \succ_{\mathbf{s}'} g$ by the assumption specifying $f, g$. Hence, by Consensus-wise Pareto, $f \succ^0_s g$. So $\succ^0_s \supseteq \bigcap_{\mathbf{s} \in X : s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$, and hence there is equality, as required.

□

For each $s \in G(X)$, since $c^0(s)$ represents $\succ^0_s$ by the confidence representation (Hill, 2016), it follows that, up to convex closure, $c^0(s) = \bigcup_{\mathbf{s} \in X : s \geqslant G(\mathbf{s})} \bigcap_i c^i(s_i)$.

By the nestedness of confidence rankings (i.e. the fact that $c$ is increasing in $o$), we have that, for any $\mathbf{s}, \mathbf{s}' \in X$, if $\mathbf{s}' \geqslant \mathbf{s}$, then $G(\mathbf{s}') \geqslant G(\mathbf{s})$, so $G$ is monotonic. Moreover, if $\succ_{\mathbf{s}} = \succ_{\mathbf{t}}$, then $G(\mathbf{s}) = G(\mathbf{t})$, so $G$ generates a well-defined function on the equivalence classes of $\mathcal{S}^n$ under the relation setting $\mathbf{s}$ and $\mathbf{t}$ equivalent if and only if $\succ_{\mathbf{s}} = \succ_{\mathbf{t}}$, which we also call $G$. So $\otimes$, defined by

$$\otimes(o_1, \ldots, o_n) = \begin{cases} \zeta \circ G(\zeta^{-1}(o_1), \ldots, \zeta^{-1}(o_n)) & (\zeta^{-1}(o_1), \ldots, \zeta^{-1}(o_n)) \in X \\ \zeta(\inf(G(X))) & \text{otherwise} \end{cases}$$

is well-defined; i.e. even if $(\zeta^{-1}(o_1), \ldots, \zeta^{-1}(o_n))$ is multi-valued, for any $\mathbf{s}, \mathbf{t} \in (\zeta^{-1}(o_1), \ldots, \zeta^{-1}(o_n))$, $\succ_{\mathbf{s}} = \succ_{\mathbf{t}}$ by the confidence decision model, and so $G((\zeta^{-1}(o_1), \ldots, \zeta^{-1}(o_n)))$ is well-defined ($G(\mathbf{s}) = G(\mathbf{t})$). Moreover, $\otimes$ is monotonic, and thus a confidence level aggregator. It follows from Claim D.3 that (1) holds up to convex closure for all $o$ with $\bigcup_{\mathbf{o} : \otimes \mathbf{o} \leqslant o} \bigcap_i c^i(o_i) \neq \varnothing$. For any $s \notin G(X)$, since $X$ is clearly upwards closed and $G$ is increasing, $s \leqslant \inf G(X)$; hence, by the nestedness of confidence rankings, $\succ^0_s \supseteq \bigcup_{s' \in G(X)} \succ^0_{s'}$. However, by Consensus-based preferences, if $f \succ^0_s g$, then $f \succ^0_{s'} g$ for some $s' \in G(X)$, so $\succ^0_s = \bigcup_{s' \in G(X)} \succ^0_{s'}$. Hence, for any $o$ with $\bigcup_{\mathbf{o} : \otimes \mathbf{o} \leqslant o} \bigcap_i c^i(o_i) = \varnothing$, $c(o) = \bigcap_{s' \in G(X)} c^0(s')$ up to convex closure, so $c^0$ is consensus preserving, as required. This establishes the Theorem.

Moreover, note that since $\otimes$ is monotonic on the domain where $\zeta^{-1}(\mathbf{o}) \in X$, any monotonic operator coinciding with $\otimes$ on this domain is also a confidence level aggregator, and represents aggregated preferences according to (1), hence establishing the 'only if' direction.

The 'if' direction is a direct consequence of (1) and Claims D.1 and

D.2.

Finally, suppose that $\otimes' \neq \otimes$ is another confidence level aggregator such that, up to convex closure, $c^0$ is a consensus-preserving aggregation of $(c^1, \ldots, c^n)$ under $\otimes'$. Let $G'(\mathbf{s}) = \zeta^{-1} \circ \otimes'(\zeta \mathbf{s})$. By the confidence representation and the fact that $c^0$ is a consensus-preserving aggregation of $(c^1, \ldots, c^n)$ under $\otimes'$, for every $\mathbf{s} \in X$, $\succ^0_{G'(\mathbf{s})} \subseteq \succ_{\mathbf{s}}$. It follows from the definition of $G$ that $G(\mathbf{s}) \leqslant G'(\mathbf{s})$ for all $\mathbf{s} \in X$. So, either $\otimes'$ coincides with $\otimes$ on $X$, or there $\mathbf{s} \in X$ with $\otimes' \mathbf{o} \neq \otimes \mathbf{o}$, so $\otimes' \mathbf{o} > \otimes \mathbf{o}$. Hence $\otimes$ is the unique $\otimes$ taking minimal values on all consensuses, as required.

*Remark* 1. Note that the use of profiles of confidence levels with $\otimes \mathbf{o}$ less than or equal to $o$, rather than just equal, as in (B.1), is a result of the general framework adopted for this result. More specifically, it is clear to see that one can prove, using arguments along the lines above, that one can replace the less than or equal with equality under the condition that, if $\otimes \mathbf{o} < o$, then there exists $\mathbf{o}' \geqslant \mathbf{o}$ with $\otimes \mathbf{o}' = o$. For an example where this condition is not satisfied, consider $O = \{a, b, c\}$ with $a > b > c$, and two agents $1, 2$. Consider $\otimes$ giving the value $c$ on $(c, c)$ and the value $a$ otherwise. Clearly, the condition is not satisfied for $b$—in fact, there is no $\mathbf{o}$ with $\otimes \mathbf{o} = b$. So there is no $\otimes$ such that (2) holds with equality in the place of the inequality.

### D.1.2 Proof of Theorem B.1

*Proof for Table 6, row 1 (affine aggregation).* Let $X$ be as defined in the proof of Theorem 1. Let

$$C = \{(\mathbf{s}, s) \in \mathcal{S}^{n+1} : \mathbf{s} \in X, \ \succ^0_s \subseteq \succ_{\mathbf{s}}\}$$
$$K = \{(\mathbf{s}, s) \in C : \succ_{\mathbf{s}} \supseteq \bigcap_{\mathbf{s}' : (\mathbf{s}', s) \in C, \ \mathbf{s}' \not\geqslant \mathbf{s}} \succ_{\mathbf{s}'}\}$$

$C$ is the set of consensuses and $K$ is the set of 'covered' consensuses—i.e. where there is consensus because the other consensuses at this $s$ 'cover' this one. For a tuple of stakes levels $\mathbf{s}$ and a stakes level $s'$, $s'_i\mathbf{s}$ is the tuple obtained by replacing the $i$th stakes level in $\mathbf{s}$ by $s'$. An individual $i$ is *non-null* if there exist $\mathbf{s}, s'_i\mathbf{s} \in X$ and $t \in \mathcal{S}$ with $(\mathbf{s}, t) \in C \backslash K$ but $(s'_i\mathbf{s}, t) \notin C$. Let $NN = \{i \in \{1, \ldots, n\} : i \text{ non-null}\}$ and $Y = \mathcal{S}^{NN} \subseteq \mathbb{R}^n$

be the subspace of $\mathcal{S}^n$ containing the stakes levels for non-null individuals only; we use $X_Y, C_Y, K_Y$ etc to refer to the projection of $X, C, K$ etc onto $Y$, $Y \times \mathbb{R}$ etc.

Define

$$
\begin{aligned}
L &= \{(\mathbf{s}, s) \in Y \times \mathbb{R} : \mathbf{s} \in X_Y,\ \succ^0_s \not\subseteq \succ_{\mathbf{s}}\} = (X_Y \times \mathbb{R}) \backslash C_Y \\
U &= \{(\mathbf{s}, s) \in C_Y : \exists s' \leqslant s, (\mathbf{s}, s') \in C_Y \backslash K_Y\}
\end{aligned}
$$

For a set $Z$, let $conv(Z)$ be the convex hull of $Z$. Note that $L, U \subseteq X_Y \times \mathbb{R}$, so $conv(L), conv(U) \subseteq conv(X_Y) \times \mathbb{R}$.

**Claim D.4.** *$conv(L) \cap conv(U) = \varnothing$.*

*Proof.* For reductio, suppose that there exist $(\mathbf{s}_1, s_1), \ldots, (\mathbf{s}_l, s_l) \in L$, $(\mathbf{t}_1, t_1), \ldots, (\mathbf{t}_m, t_m) \in U$, $\alpha_1, \ldots, \alpha_l, \beta_1, \ldots, \beta_m \in [0, 1]$ with $\sum_{i=1}^{l} \alpha_i = \sum_{i=1}^{m} \beta_i = 1$, $\sum_{i=1}^{l} \alpha_i \mathbf{s}_i = \sum_{i=1}^{m} \beta_i \mathbf{t}_i$ and $\sum_{i=1}^{l} \alpha_i s_i = \sum_{i=1}^{m} \beta_i t_i$. Without loss of generality, the $t_i$ can be chosen to be minimal such that $(\mathbf{t}_i, t_i) \in U$. It follows from Consensus Independence (extending to tuples to take any value off $NN$ for which there is consensus, if necessary) that $\sum_{i=1}^{l} \alpha_i s_i < \sum_{i=1}^{m} \beta_i t_i$, which is a contradiction.

□

**Claim D.5.** *$conv(L)$ is open in the subspace topology on $conv(X_Y) \times \mathbb{R}$.*

*Proof.* Note that $L^c \cap (X_Y \times \mathbb{R}) = C_Y = \{(\mathbf{s}, s) \in Y \times \mathbb{R} : \mathbf{s} \in X,\ G(\mathbf{s}) \geqslant s\}$, where $G$ is as defined in the proof of Theorem 1. Since $G$ is increasing, by Proposition D.2 and the nestedness of the preferences orders at different stakes levels, $L^c \cap (X_Y \times \mathbb{R})$ is closed. Hence $L$ is open in the subspace topology on $X_Y \times \mathbb{R}$. Given the finite-dimensationality of $X_Y \times \mathbb{R}$, it follows that the convex hull $conv(L)$ is open in the subspace topology on $conv(X_Y) \times \mathbb{R}$. □

By the previous claims and a separating hyperplane theorem (Rockafellar, 1970, Thm 11.3), there exists a linear function $\phi : \mathbb{R}^{NN} \to \mathbb{R}$ and $\chi \in \mathbb{R}$ with $\phi((\mathbf{s}, s)) < \chi$ for all $(\mathbf{s}, s) \in conv(L)$, and $\phi((\mathbf{s}, s)) \geqslant \chi$ for all $(\mathbf{s}, s) \in conv(U)$. Since it is linear and $C_Y \neq \varnothing$, without loss of generality, $\phi, \chi$ can be chosen so there exist $w_i$, $i \in NN$ such that $\phi((\mathbf{s}, s)) = s - \sum_i w_i s_i$. Define $G_{aff} : \mathbb{R}^n \to S$ by $G_{aff}(\mathbf{s}) = \sum_{i \in NN} w_i s_i + \chi$. Note that $G_{aff}$ is an affine function on $\mathbb{R}^n$, with zero weights on $i \notin NN$. By construction, $s < G_{aff}(\mathbf{s})$ for all $(\mathbf{s}, s) \in conv(L)$, and $s \geqslant G_{aff}(\mathbf{s})$ for all $(\mathbf{s}, s) \in conv(U)$.

We first show that $w_i > 0$ for all $i \in NN$. By the nestedness of confidence rankings, for any $\mathbf{s}', \mathbf{s} \in Y$, $\mathbf{s}' \geqslant \mathbf{s}$, if $(\mathbf{s}, s) \in L$, then $(\mathbf{s}', s) \in L$. For reductio, suppose, for some $k$, that $w_k < 0$, and consider $(\mathbf{s}, s') \in L$. By construction, $s$, with $s - \sum_i w_i s_i = \chi$, is such that $(\mathbf{s}, s) \notin L$. Consider $\mathbf{s}' = (s_1, \ldots, s_k - \frac{s-s'}{w_k}, \ldots, s_n)$. $\mathbf{s}' \geqslant \mathbf{s}$ since $w_k < 0$, so $(\mathbf{s}', s') \in L$. However, $s' - \sum_i w_i s'_i = \chi$, contradicting the established properties of $\phi$. Hence $w_i \geqslant 0$ for all $i \in NN$. Suppose now that for some $i \in NN$, $w_i = 0$. By the nestedness of the confidence representation and the definition of $NN$, there exists $\mathbf{s} \in X$, $s', t$ such that $(\mathbf{s}, t) \in U$ and $(s'_i\mathbf{s}, t) \in L$; however, since $w_i = 0$, $G_{aff}(\mathbf{s}) = G_{aff}(s'_i\mathbf{s})$, which contradicts the definition of $G_{aff}$. So, for all $i \in NN$, $w_i \neq 0$. Hence $w_i > 0$ for all $i \in NN$, and $G_{aff}$ is monotonic.

**Claim D.6.** For all $s \in G_{aff}(X)$, $\succ^0_s$ is represented by $\bigcup_{\mathbf{s} \in X : s = G_{aff}(\mathbf{s})} \bigcap_i c^i(s_i)$ in the Bewley sense: for all $f, g \in \mathcal{A}$, $f \succ^0_s g$ if and only if:

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \quad \text{for all } p \in \bigcup_{\mathbf{s} \in X : s = G_{aff}(\mathbf{s})} \bigcap_i c^i(s_i) \tag{D.4}$$

*Proof.* Fix a stakes level $s$, with $s \geqslant \inf G_{aff}(X)$. For any $\mathbf{s} \in X$ with $G_{aff}(\mathbf{s}) = s$, by the construction of $\phi$ and the definition of $NN$, $\succ^0_s \subseteq \succ_{\mathbf{s}}$. So $\succ^0_s \subseteq \bigcap_{\mathbf{s} \in X : s = G_{aff}(\mathbf{s})} \succ_{\mathbf{s}}$.

We now establish the opposite containment. By Consensus-wise Pareto, $\succ^0_s \supseteq \bigcap_{\mathbf{s} : (\mathbf{s}, s) \in C} \succ_{\mathbf{s}}$. Consider any $\mathbf{s}'$ such that $\succ^0$ respects the consensus $\succ_{\mathbf{s}'}$ at $s$—so $(\mathbf{s}', s) \in C$—and $G_{aff}(\mathbf{s}') < s$. Then by the fact that the $w_i \geqslant 0$ for all $i$, there exists $\mathbf{s} \geqslant \mathbf{s}'$ with $G_{aff}(\mathbf{s}) = s$; by the nestedness of confidence rankings and the preference representation, $\succ_{\mathbf{s}'} \supseteq \succ_{\mathbf{s}} \supseteq \bigcap_{\mathbf{s} : (\mathbf{s}, s) \in C,\ G_{aff}(\mathbf{s}) \geqslant s} \succ_{\mathbf{s}}$. Since this holds for all such $\mathbf{s}'$, $\succ^0_s \supseteq \bigcap_{\mathbf{s} : (\mathbf{s}, s) \in C,\ G_{aff}(\mathbf{s}) \geqslant s} \succ_{\mathbf{s}} = \bigcap_{\mathbf{s} : G_{aff}(\mathbf{s}) = s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s} : (\mathbf{s}, s) \in C,\ G_{aff}(\mathbf{s}) > s} \succ_{\mathbf{s}}$, where the equality is due to the construction of $G_{aff}$. Now consider any $\mathbf{s}'$ with $(\mathbf{s}', s) \in C$ and $G_{aff}(\mathbf{s}') > s$. If $(\mathbf{s}', s) \notin K$, then $(\mathbf{s}', s) \in U$, contradicting the fact that $G_{aff}(\mathbf{s}') > s$ and the construction of $G_{aff}$. Hence $(\mathbf{s}', s) \in K$, so $\succ_{\mathbf{s}'} \supseteq \bigcap_{\mathbf{s}'':(\mathbf{s}'',s)\in C,\ \mathbf{s}''\not\geq\mathbf{s}'} \succ_{\mathbf{s}''}$. So $\bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s} \succ_{\mathbf{s}} = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s, \mathbf{s}\not\geq\mathbf{s}'} \succ_{\mathbf{s}}$. Since this holds for all such $\mathbf{s}'$, it follows that $\bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s} \succ_{\mathbf{s}} = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}}$, so $\succ^0_s \supseteq \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}}$.

So $\succ^0_s = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}}$; it follows from Claim D.2 that (D.4) holds for all $s \in G_{aff}(X)$.

□

Since $c^0(s)$ represents $\succ^0_s$ by the confidence representation (Hill, 2016), it follows that, up to convex closure, $c^0(s) = \bigcup_{\mathbf{s}\in X:G_{aff}(\mathbf{s})=s} \bigcap_i c^i(s_i)$.

Define $\otimes$ by

$$\otimes\mathbf{o} = \begin{cases} \sum w_i o_i + \chi & \sum w_i o_i + \chi \geqslant 0 = \min O \\ 0 & \text{otherwise} \end{cases}$$

Clearly, this is an affine confidence level aggregator. Moreover, by Claim D.6 and the fact that $\zeta$ is the identity, (B.1) holds up to convex closure for every $o$ with $\bigcup_{\mathbf{o}:\otimes o_i=o} \bigcap_i c^i(o_i) \neq \varnothing$. By a similar argument to that used in the proof of Theorem 1, the representation extends to other $o \in O$ as required. Hence, up to convex closure, $c^0$ is a consensus preserving with affine aggregator $\otimes$ as required.

For the necessity of the Consensus Independence axiom, suppose that there is an affine aggregator $\otimes$ representing preferences. Consider any $\mathbf{s_1}, \ldots, \mathbf{s_l}, \mathbf{t_1}, \ldots, \mathbf{t_m}$ exhibiting consensus, and $\alpha_1, \ldots, \alpha_l, \beta_1, \ldots, \beta_m \in [0,1]$ with $\sum_{k=1}^l \alpha_k = \sum_{k=1}^m \beta_k = 1$ and $\sum_{k=1}^l \alpha_k \mathbf{s_k} = \sum_{k=1}^m \beta_k \mathbf{t_k}$. If $\succ^0$ does not respect the consensuses $\succ_{\mathbf{s_k}}$ at $s_k$, then $c^0(s_k) \not\supseteq \bigcap_i c^i((\mathbf{s_k})_i)$, whereas, by the aggregation rule $c^0(\sum_i w_i(\mathbf{s_k})_i + \chi) \supseteq \bigcap_i c^i((\mathbf{s_k})_i)$, so, by the nestedness of confidence rankings, $s_k < \sum_i w_i(\mathbf{s_k})_i + \chi$. If this holds for all $k$, then $\sum_{k=1}^l \alpha_k s_k < \sum_i w_i \sum_{k=1}^l \alpha_k(\mathbf{s_k})_i + \chi$. Similarly, if $\succ_{\mathbf{t_k}}$ is an uncovered consensus at $t_k$ then, by the confidence representation of preferences, $c^0(t_k) \supseteq \bigcap_i c^i((\mathbf{t_k})_i)$ and $\bigcap_i c^i((\mathbf{t_k})_i) \nsubseteq \bigcup_{\mathbf{s}\not\geq\mathbf{t_k},\ (\mathbf{s},t_k)\in C} \bigcap_i c^i(s_i) \subseteq c^0(t_k)$. By the aggregation representation, it follows that $c^0(\sum_i w_i(\mathbf{t_k})_i + \chi) = \bigcup_{\mathbf{s}:\ \sum_i w_i s_i = \sum_i w_i(\mathbf{t_k})_i} \bigcap_i c^i(s_i) \subseteq \bigcap_i c^i((\mathbf{t_k})_i) \cup \bigcup_{\mathbf{s}\not\geq\mathbf{t_k},\ (\mathbf{s},t_k)\in C} \bigcap_i c^i(s_i) \subseteq c^0(t_k)$, so, by the nestedness of confidence rankings, $\sum_i w_i(\mathbf{t_k})_i + \chi \leqslant t_k$. So if $\succ_{\mathbf{t_k}}$ is an uncovered consensus at $t_k$ for each $k$, it follows that $\sum_{k=1}^m \beta_k t_k \geqslant \sum_i w_i \sum_{k=1}^m \beta_k(\mathbf{t_k})_i + \chi$. Since, $\sum_i w_i \sum_{k=1}^l \alpha_k(\mathbf{s_k})_i = \sum_i w_i \sum_{k=1}^m \beta_k(\mathbf{t_k})_i$, it follows that $\sum_{k=1}^{m} \beta_k t_k > \sum_{k=1}^{l} \alpha_k s_k$, as required.

□

*Proof for Table 6, row 2 (generalised average aggregation).* We show that there exists a representation of the sort obtained in the proof of part i. where the weights are equal. For reductio, suppose not, and consider a representation with an affine aggregator with $w_j > w_k$ for some $j, k$. First, by Neutrality and Non-degeneracy, $NN = \{1, \ldots, n\}$, so $w_j, w_k \neq 0$.

First consider the case where there exists $s$ and $\mathbf{s}$ such that $(\mathbf{s}, s) \in C$, $\mathbf{s}$ is a maximum, under $\geqslant$, of $\{\mathbf{s}' : (\mathbf{s}', s) \in C\}$, and $s_j \neq s_k$; take any such $s$ and $\mathbf{s}$. By the upper semicontinuity of confidence rankings, for any strictly decreasing sequences $t_l \to s_j$ and $t'_l \to s_k$, $\bigcap_{i \neq j} c^i(s_i) \cap c^k(t'_l) \to \bigcap_i c^i(s_i)$ and $\bigcap_{i \neq k} c^i(s_i) \cap c^j(t_l) \to \bigcap_i c^i(s_i)$ as $l \to \infty$. By the fact that $\mathbf{s}$ is a maximum, $((t_l)_j\mathbf{s}, s) \notin C$, $((t'_l)_k\mathbf{s}, s) \notin C$ for all $l$. By the affine aggregator representation and the upper semicontinuity of confidence rankings, for each $s'' > s$, there exist $m_t, m_{t'}$ with $((t_l)_j\mathbf{s}, s'') \in C$ and $((t'_l)_k\mathbf{s}, s'') \in C$ for all $l > m_t$ and $l > m_{t'}$. In particular $G_{aff}(t_j\mathbf{s}) > s$ and $G_{aff}(t'_k\mathbf{s}) > s$ for all $t > s_j$, $t' > s_k$, where $G_{aff}$ is as in the proof of part i., though $G_{aff}((t_l)_j\mathbf{s}) \to G_{aff}(\mathbf{s})$ and $G_{aff}((t'_l)_k\mathbf{s}) \to G_{aff}(\mathbf{s})$ as $l \to \infty$, so by the continuity of the affine representation, $G_{aff}(\mathbf{s}) = s$.

If $s_j > s_k$, then $G_{aff}((s_k)_j(s_j)_k\mathbf{s}) < s$, by the form of $G_{aff}$, the fact that $w_j > w_k$ and the rearrangement inequality. Hence, by the continuity of the representation, for some $t > s_k$, $G_{aff}(t_j(s_j)_k\mathbf{s}) < s$, from which it follows that $(t_j(s_j)_k\mathbf{s}, s) \in C$. Since $t_j(s_j)_k\mathbf{s}$ is a permutation of $t_k\mathbf{s}$, it follows from Neutrality that $(t_k\mathbf{s}, s) \in C$, contradicting the maximality of $\mathbf{s}$. If $s_j < s_k$, then $G_{aff}((s_k)_j(s_j)_k\mathbf{s}) > s$. By the construction of $\mathbf{s}$ there exists $s'' < G_{aff}((s_k)_j(s_j)_k\mathbf{s})$ and $t > s_k$ with $(t_k\mathbf{s}, s'') \in C$. Moreover, since $t$ can be chosen such that there exists $t' > t$ with $(t'_k\mathbf{s}, s'') \notin C$, $t$ can be chosen so that $(t_k\mathbf{s}, s'') \notin K$. By Neutrality, it follows that $(t_j(s_j)_k\mathbf{s}, s'') \in C\backslash K$, so $(t_j(s_j)_k\mathbf{s}, s'') \in U$, contradicting the construction of $G_{aff}$ and the fact that $G_{aff}(t_j(s_j)_k\mathbf{s}') \geqslant G_{aff}((s_k)_j(s_j)_k\mathbf{s}') > s''$.

Now consider the case where there does not exist $s$ and $\mathbf{s}$ such that $(\mathbf{s}, s) \in C$, $\mathbf{s}$ is a maximum, under $\geqslant$, of $\{\mathbf{s}' : (\mathbf{s}', s) \in C\}$, and $s_j \neq s_k$. Hence, for all $s$ and $\mathbf{s}$ such that $(\mathbf{s}, s) \in C$ and $\mathbf{s}$ is a maximum, under $\geqslant$, of $\{\mathbf{s}' : (\mathbf{s}', s) \in C\}$, $s_j = s_k$. For any $\mathbf{s}$, let $\hat{\mathbf{s}}$ be such that: $\hat{s}_i = s_i$

when $i \neq j, k$, $\hat{s}_j = \hat{s}_k = \max\{s_j, s_k\}$. So, in the case under consideration, for every $\mathbf{s}$ with $s_j \neq s_k$ and every stakes level $s$, $(\mathbf{s}, s) \in C$ if and only if $(\hat{\mathbf{s}}, s) \in C$.

Hence the map $\psi : \mathcal{S}^n \to \mathcal{S}^{n-1}$, defined by $\psi(\mathbf{s})_i = s_i$ for $i \neq j, k$ and $\psi(\mathbf{s})_j = \max\{s_j, s_k\}$, is a well-defined map sending $C$ to $\psi(C) = \hat{C}$ which is such that $\psi^{-1}(\hat{C}) = C$. Hence images of other sets in the proof of part i., which are defined in terms of $C$, can be defined in terms of $\hat{C}$ and have the same pull-back property. It follows that the argument in the proof of part i. goes through, yielding a representation of $c^0$ in terms of an affine function $\widehat{G_{aff}} : \mathbb{R}^{n-1} \to \mathbb{R}$ of the following form: for all $s \geqslant \inf \widehat{G_{aff}}(\hat{X})$:

$$c^0(s) = \bigcup_{\mathbf{s} \in \hat{X} : s = \widehat{G_{aff}}(\mathbf{s})} \bigcap_{i \neq k} c^i(s_i) \cap c^k(s_j)$$

up to convex closure. Letting $\widehat{G_{aff}}(\mathbf{s}) = \sum_{i \neq k} w_i s_i + \chi$, define $G'_{aff} : \mathbb{R}^n \to \mathbb{R}$ by $\widehat{G_{aff}}(\mathbf{s}) = \sum_{i \neq j,k} w_i s_i + \frac{w_j}{2} s_j + \frac{w_j}{2} s_k + \chi$. Noting that, for all $\mathbf{s} \in \mathcal{S}^n$ with $s_j = s_k$, $G'_{aff}(\mathbf{s}) = \widehat{G_{aff}}(\mathbf{s}|_{\{1,\ldots n\} \setminus \{k\}})$, we have that, for all $s \geqslant \inf G'_{aff}(X)$,

$$c^0(s) = \bigcup_{\mathbf{s} \in X : s = G'_{aff}(\mathbf{s}),\ s_j = s_k} \bigcap_{i=1}^{n} c^i(s_i)$$

up to convex closure.

For any $\mathbf{s} \in X$ with $s_j < s_k$ and $(\mathbf{s}, s) \in C$, it follows that $(\hat{\mathbf{s}}, s) \in C$ by the specification of the case. Since $\bigcap_i c^i(s_i) \subseteq \bigcap_i c^i(\hat{s}_i) \subseteq \bigcup_{\mathbf{s}' \in X : s = G'_{aff}(\mathbf{s}'),\ s_j = s_k} \bigcap_i c^i(s_i)$, it follows that, for all $s$ with $(\mathbf{s}, s) \in C$, $c^0(s) = \bigcup_{\mathbf{s} \in X : s = G'_{aff}(\mathbf{s})} \bigcap_{i=1}^{n} c^i(s_i)$, up to convex closure. By the specification of the case, this implies that this holds for all $s \in G'_{aff}(X)$. So there is an affine aggregator representation with equal weights for $j$ and $k$, as required.

Necessity of Neutrality is straightforward.

□

*Proof for Table 6, row 3 (non-neutral maximum aggregator).* Consider $G$ as defined in the proof of Theorem 1. By Consensus Join, for any $\mathbf{s}_k$, $k \in K$, $G(\bigvee_{k \in K} \mathbf{s}_k) \leqslant \sup_{k \in K} G(\mathbf{s}_k)$. However, by the monotonicity of $G$, since $\bigvee_{k \in K} \mathbf{s}_k \geqslant \mathbf{s}_k$ for all $k \in K$, $G(\bigvee_{k \in K} \mathbf{s}_k) \geqslant \sup_{k \in K} G(\mathbf{s}_k)$, so $G(\bigvee_{k \in K} \mathbf{s}_k) = \sup_{k \in K} G(\mathbf{s}_k)$. For each $s \in G(X)$, consider $\mathbf{t}^s = \bigvee_{\mathbf{s} : G(\mathbf{s}) \leqslant s} \mathbf{s}$. By the previous observation, $G(\mathbf{t}^s) = s$ and for any $\mathbf{s}$ with $s_i > t_i^s$ for some $i$, $G(\mathbf{s}) > s$. Since, for any $\mathbf{s}$, if $\mathbf{s} \leqslant \mathbf{t}^s$, then $G(\mathbf{s}) \leqslant s$ by the monotonicity of $G$, we have that, for all $\mathbf{s}$, $G(\mathbf{s}) > s$ if and only if there exists $i$ with $s_i > t_i^s$. Hence $G(\mathbf{s}) < s$ if and only if there exists $s' < s$ with $s_i \leqslant t_i^{s'}$ for all $i$. Hence $G(\mathbf{s}) = s$ if and only if $\mathbf{s} \leqslant \mathbf{t}^s$ and there is no $s' < s$ with $\mathbf{s} \leqslant \mathbf{t}^{\mathbf{s}'}$.

Moreover, since, by the nestedness of the confidence representation, $\bigcap_i c^i(s_i) \subseteq \bigcap_i c^i(t_i^s)$ for all $\mathbf{s}$ with $G(\mathbf{s}) \leqslant s$, it follows that $\bigcap_i c^i(t_i^s) = \bigcup_{\mathbf{s}\in X:s=G(\mathbf{s})} \bigcap_i c^i(s_i) = \bigcup_{\mathbf{s}\in X:s\geqslant G(\mathbf{s})} \bigcap_i c^i(s_i)$. So, up to convex closure, $c^0(s) = \bigcap_i c^i(t_i^s)$.

For $i = 1, \ldots, n$, define $\psi_i : O \to O$ by $\psi_i(o) = \zeta(\min\{s \in \mathcal{S} : \zeta(t_i^s) \geqslant o\})$. Since $\zeta$ is strictly increasing and, by the confidence representation, $t_i^s$ is increasing in $s$ for all $i$, $\psi_i$ is increasing for all $i$. For any $\mathbf{o} \in O^n$, $s \in \mathcal{S}$ and $\mathbf{s} \in \zeta^{-1}(\mathbf{o})$, $G(\mathbf{s}) = s$ if and only if $\mathbf{s} \leqslant \mathbf{t}^{\mathbf{s}}$ and $\mathbf{s} \nleqslant \mathbf{t}^{\mathbf{s}'}$ for all $s' < s$, which is the case if and only if $\max_i \psi_i(o_i) = s$. Hence $G(\zeta^{-1}(\mathbf{o}))$ is well-defined, and $G(\zeta^{-1}(\mathbf{o})) = \max_i \psi_i(o_i)$. Defining $\otimes$ by $\otimes\mathbf{o} = \max_i \psi_i(o_i)$, we thus have that, for every $o$ with $\bigcup_{\mathbf{o}:\otimes o_i=o} \bigcap_i c^i(o_i) \neq \varnothing$, (B.1) holds with $\otimes$, up to convex closure. By a similar argument to that used in the proof of Theorem 1, the representation extends to other $o \in O$ as required. Since the $\psi_i$ are increasing, $\otimes$ is monotonic, and hence a generalized maximum aggregator. Hence, up to convex closure, $c^0$ is a consensus preserving with generalised maximum aggregator $\otimes$ as required.

The proof of necessity of Consensus Join is straightforward.

□

*Proof for Table 6, row 4 & 5 (generalised maximum & minimum aggregators).* We present the proof for the maximum aggregator; the case of the minimum aggregator is similar. Consider $\mathbf{t}^{\mathbf{s}}$, as defined in the proof of part iii; we show that $t_j^s = t_k^s$ for all $j, k$. For reductio, suppose that this is not the case for some $j, k$, and suppose without loss of generality that $t_j^s > t_k^s$. By Neutrality, $G((t_k^s)_j(t_j^s)_k\mathbf{t}^{\mathbf{s}}) = G(\mathbf{t}^{\mathbf{s}}) = s$; but since $t_j^s > t_k^s$, it follows by the properties of $G$ established in the proof of part iii. that $G((t_k^s)_j(t_j^s)_k\mathbf{t}^{\mathbf{s}}) > G(\mathbf{t}^{\mathbf{s}}) = s$, which is a contradiction. So $t_j^s = t_k^s$ for all $j, k$ and $s$. Hence, for $\psi_i$ as defined in the proof of part iv., $\psi_j(o) = \psi_k(o) = \psi(o)$ for all $j, k$ and $o \in O$, whence $\otimes$ as defined in that proof of the proof can be written as $\otimes\mathbf{o} = \max_i \psi(o_i) = \psi(\max_i o_i)$. Hence it is a maximum aggregator, as required.

□

## D.2 Proofs of results in Section 3

For the results in this Section, it is useful to recall the definition of $f$-divergence (Liese and Vajda, 1987). Given a convex continuous function $f : (0, \infty) \to \mathbb{R}$ satisfying $f(1) = 0$, the $f$-divergence of $p \in \Delta$ with respect to $q \in \Delta$ is defined by

$$D_f(p\|q) = \begin{cases} \int_\Omega f\left(\frac{dp}{dq}(\omega)\right) dq(\omega) & p \text{ absolutely continuous with respect to } q \\ \infty & \text{otherwise} \end{cases} \tag{D.5}$$

where $\frac{dp}{dq}$ is the Radon–Nikodym derivative.

*Proof of Proposition 1.* By (2), the centre of $c$ is:

$$\begin{aligned} \arg\min_{p\in\Delta} \otimes(\iota^1(p), \ldots, \iota^n(p)) &= \arg\min_{p\in\Delta}(\sum_{i=1}^{n} \frac{1}{n}\iota^i(p) + \chi) \\ &= \arg\min_{p\in\Delta} \sum_{i=1}^{n} \frac{1}{n}\iota^i(p) \end{aligned}$$

For the first row of Table 4, $\iota^i(p) = w^i \sum_{\omega\in\Omega}(p(\omega) - p_i(\omega))^2$, so the centre of $c$ is $p = \arg\min_{p\in\Delta} \sum_{i=1}^{n} w^i \sum_{\omega\in\Omega}(p(\omega) - p_i(\omega))^2$. It is well-known that this is the mean of the distributions: the FOC is $\frac{d}{dp(\omega)} = 2\sum_{i=1}^{n} w^i(p(\omega) - p_i(\omega)) = 0$ for each $\omega \in \Omega$, yielding $p(\omega) = \sum_{i=1}^{n} \frac{w^i}{\sum_{i=1}^{n} w^i} p_i(\omega)$ for every $\omega \in \Omega$, which belongs to $\Delta$.

The second and fourth rows of the Table involve generating distances that are $f$-divergences. They will be consequences of the following more general claim.

**Claim D.7.** Let $D_f$ be an $f$-divergence generated by a strictly convex differentiable function $f : (0, \infty) \to \mathbb{R}$ satisfying $f(1) = 0$, and suppose that

$\bigcap_{i:w_i>0} \operatorname{supp} p_i \neq \varnothing$. Let $c$ be the confidence aggregation under the average confidence-level aggregator of the $w^i$ $D_f$-confidence rankings generated by $p_i$. Then its centre is $p \in \Delta$ such that there exists $\lambda \in \mathbb{R}$ with

$$\sum_{i=1}^{n} w_i f'\left(\frac{dp}{dp_i}(\omega)\right) = \lambda \qquad \text{for } \mu\text{-almost every } \omega \in \Omega \tag{D.6}$$

where $\mu \in \Delta$ is defined by $\mu = \sum_{i=1}^{n} \frac{w^i}{\sum_{i=1}^{n} w^i} p_i$.

*Proof.* As noted above, the centre of $c$ is

$$\arg\min_{q\in\Delta} \sum_{i=1}^{n} w_i D_f(q|p_i) = \arg\min_{q\in\Delta} \sum_{i=1}^{n} w_i \int_{\Omega} f\left(\frac{dq}{dp_i}\right) dp_i. \tag{D.7}$$

Since there is non-empty intersection of the supports of the $p_i$ with positive weight, there exists some probability measure that is absolutely continuous with respect to all these $p_i$. Hence the set of optima (the centre) is not the whole space.

Moreover, since $f$ is strictly convex, each map $q \mapsto D_f(q|p_i)$ is convex, and the weighted sum is strictly convex whenever at least one weight is strictly positive. Therefore any interior optimum is unique.

To derive the first-order condition, first note that all $p_i$ for which $w_i > 0$ are clearly absolutely continuous with respect to $\mu$. Since any optimum must be absolutely continuous with respect to all $p_i$ with $w_i > 0$, it must be absolutely continuous with respect to $\mu$; henceforth we limit consideration to such $q$. The Lagrangian can thus be written as:

$$L(q) = \sum_{i=1}^{n} w_i \int_{\Omega} \frac{dp_i}{d\mu} f\left(\frac{\frac{dq}{d\mu}}{\frac{dp_i}{d\mu}}\right) d\mu - \lambda\left(\int_{\Omega} \frac{dq}{d\mu} d\mu - 1\right)$$

Let $h$ be any perturbation satisfying

$$\int_{\Omega} h d\mu = 0.$$

Then

$$\frac{d}{d\varepsilon}L(q+\varepsilon h,\lambda)\Big|_{\varepsilon=0}=\int_{\Omega}\left[\sum_{i=1}^{n}w_i f'\left(\frac{\frac{dq}{d\mu}}{\frac{dp_i}{d\mu}}\right)-\lambda\right]hd\mu.$$

At an (interior) optimum $p$ this derivative must vanish for every admissible $h$. By the fundamental lemma of the calculus of variations, this holds if and only if:

$$\sum_{i=1}^{n}w_i f'\left(\frac{dp}{dp_i}(\omega)\right)=\lambda \qquad \text{for } \mu\text{-almost every } \omega\in\Omega.$$

It follows from the fact that $f'$ is strictly increasing that any 'boundary' $q$ with $q(E)=0$ for some $E$ such that $p_i(E)>0$ for all $i$ with $w_i>0$ yields a higher value in the optimatisation problem (D.7) than some 'interior' $q$ with $q(E)>0$ whenever $p_i(E)>0$ for all $i$ with $w_i>0$. Hence the optimal $p$ in (D.7) is 'interior', and thus satisfies this condition. Thus the centre of $c$ is the unique probability measure $p\in\Delta$ satisfying this condition.

□

For the second row of the Table, note that relative entropy is a $f$-divergence with $f(t)=t\log t$. This is strictly convex and differentiable, so Claim D.7 applies. First note that, if the intersection of the supports of the probability measures with positive weight is empty, then the geometric pool is not defined; we may thus focus on the case of non-empty intersection. Plugging the form of $f$ into the condition from the Claim yields that, for $\mu$-almost every $\omega\in\Omega$:

$$\sum_{i=1}^{n}w_i\left(\log\frac{dp}{dp_i}(\omega)+1\right)=\lambda$$

whence

$$\left(\sum_{i=1}^{n}w_i\right)\log\frac{dp}{d\mu}(\omega)=\sum_{i=1}^{n}w_i\log\frac{dp_i}{d\mu}(\omega)+\text{constant}$$

so

$$\frac{dp}{d\mu}(\omega) \propto \prod_{i=1}^{n} \left( \frac{dp_i}{d\mu}(\omega) \right)^{\frac{w_i}{\sum_{i=1}^{n} w_i}}$$

This is the definition of $p$ being the geometric pool with weights $\frac{w_i}{\sum_{i=1}^{n} w_i}$.

For the fourth row of the Table, similarly, the Gini divergence is a $f$-divergence with strictly convex and differentiable $f(t) = (t-1)^2$. First note that, if the intersection of the supports of the probability measures with positive weight is empty, then the harmonic pool is not defined; we may thus focus on the case of non-empty intersection. Plugging the form for $f$ into the condition from the Claim yields that, for $\mu$-almost every $\omega \in \Omega$:

$$\sum_{i=1}^{n} 2w_i \left( \frac{dp}{dp_i}(\omega) - 1 \right) = \lambda$$

whence

$$\frac{dp}{d\mu} \sum_{i=1}^{n} \frac{w_i}{\frac{dp_i}{d\mu}} = \text{constant}$$

so

$$\frac{dp}{d\mu} \propto \frac{1}{\sum_{i=1}^{n} \frac{w_i}{\frac{dp_i}{d\mu}}}$$

This is the definition of $p$ being a weighted harmonic pool with weights $\frac{w_i}{\sum_{i=1}^{n} w_i}$.

The third row in the Table can be partially derived from Claim D.7: since reverse relative entropy is also a $f$-divergence, it follows from the Claim, under the assumption that the supports of $p_i$ with positive weight intersect. Since the linear pool is defined even in the absence of this assumption, we provide a direct proof that does not require it.

As noted above, the centre is:

$$\arg\min_{q\in\Delta}\sum_{i=1}^{n} w_i R(p_i\|q).$$

Note that $p_i \ll \mu$ for all $i$ by construction. By the definition of relative entropy, any minimiser $p$ satisfies $p_i \ll p$ for all $i$, so we need only consider $q$ for which this holds. Note that, for any such $q$, $\mu \ll q$.

Using the chain rule for Radon–Nikodym derivatives, we have:

$$\begin{aligned} R(p_i\|q) &= \int_\Omega \log\left(\frac{dp_i}{dq}\right) dp_i \\ &= \int_\Omega \log\left(\frac{dp_i}{d\mu}\frac{d\mu}{dq}\right) dp_i \\ &= R(p_i|\mu) + \int_\Omega \log\left(\frac{d\mu}{dq}\right) dp_i. \end{aligned}$$

Hence

$$\begin{aligned} \sum_{i=1}^{n} w_i R(p_i|q) &= \sum_{i=1}^{n} w_i R(p_i|\mu) + \sum_{i=1}^{n} w_i \int_\Omega \log\left(\frac{d\mu}{dq}\right) dp_i \\ &= \sum_{i=1}^{n} w_i R(p_i|\mu) + \left(\sum_{i=1}^{n} w_i\right) \int_\Omega \log\left(\frac{d\mu}{dq}\right) d\mu \\ &= \sum_{i=1}^{n} w_i R(p_i|\mu) + \left(\sum_{i=1}^{n} w_i\right) R(\mu|q). \end{aligned}$$

The first term is constant with respect to $q$. Therefore minimizing $\sum_i w_i R(p_i|q)$ is equivalent to minimizing $R(\mu|q)$. Since relative entropy is non-negative and vanishes if and only if its two arguments coincide, the optimum $p = \mu$, which is the linear pool with weights $\frac{w_i}{\sum_{i=1}^{n} w_i}$, as required. □

The next two results and proofs adopt the notation from Example 3.1.

**Proposition D.3.** *Under the conditions and setup of Example 3.1, let* $c^{Eucl}$ *be the* $w^L$ *Euclidean confidence ranking generated by* $p^L$ *(with* $\omega' = \omega_R$*),* $c^{D_f}$ *be the* $w^L$ $D_f$ *confidence ranking generated by* $p^L$*, for an* $f$*-divergence* $D_f$*. Then, for all* $o \in O$, $c^{Eucl}(o) \subseteq \mathcal{L}_\epsilon$ *if and only if* $c^{Eucl}(o) \subseteq \mathcal{R}_\epsilon$*, and* $c^{D_f}(o) \subseteq \mathcal{L}_\epsilon$ *if and only if* $c^{D_f}(o) \subseteq \mathcal{R}_\epsilon$*.*

*Since relative entropy, reverse relative entropy and Gini are $f$-divergences, the same holds for confidence rankings generated by these distances.*

*Proof.* It suffices to show that the appropriate distance (or, equivalently $\iota$-value) between $p$ and the closest $q$ with $q(L) = 0.9 - \epsilon$ is the same as the distance to the closest $q'$ with $q'(R) = 0.1 + \epsilon$.

First consider the Euclidean case. The distance function involved is a function of $p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_{LR}), p(\omega_L), p(\omega_N)$; write this function as $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_{LR}), p(\omega_L), p(\omega_N))$. More specifically:

$$
\begin{aligned}
&\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_{LR}), p(\omega_L), p(\omega_N)) \\
=&(p(\omega_{LR}) - p^L(\omega_{LR}))^2 + (p(\omega_L) - p^L(\omega_L))^2 + ((p(\omega_N) - p^L(\omega_N))^2 \\
&+ ((1 - p(\omega_{LR}) - p(\omega_L) - p(\omega_N)) - (1 - p^L(\omega_{LR}) - p^L(\omega_L) - p^L(\omega_N)))^2
\end{aligned}
$$

Note that, since $p^L(\omega_{LR}) = p^L(\omega_N)$, $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_{LR}), p(\omega_L), p(\omega_N)) = \phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_N), p(\omega_L), p(\omega_{LR}))$ for all $p$.

Let $q$ minimise the distance from $p^L$ among all $p$ with $p(L) = 0.9 - \epsilon$. I.e. $q$ minimises $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), q(\omega_{LR}), q(\omega_L), q(\omega_N))$ among all $p$ with $p(L) = 0.9 - \epsilon$. Hence, by the previous observation, $q$ minimises $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), q(\omega_N), q(\omega_L), q(\omega_{LR}))$ among all $p$ with $p(L) = p(\omega_L) + p(\omega_{LR}) = 0.9 - \epsilon$. Define $q'$ by $q'(\omega_{LR}) = q(\omega_N)$, $q'(\omega_L) = q(\omega_L)$, $q'(\omega_N) = q(\omega_{LR})$. By the previous observation, $q'$ minimises $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), q'(\omega_{LR}), q$ among all $p$ with $p(R^c) = p(\omega_L) + p(\omega_N) = 0.9 - \epsilon$. So $q'$ minimises the distance from $p^L$ among all $p$ with $p(R) = 0.1 + \epsilon$. By the previous observation, the distance between $q$ and $p^L$ is the same as the distance between $q'$ and $p^L$, as required.

For the $f$-divergence case, the result follows from the following claim.

**Claim D.8.** Let $p \in \Delta$ and $D_f$ be an $f$-divergence for strictly convex $f : (0, \infty) \to \mathbb{R}$. For an event $E$ with $p(E) \in (0, 1)$, consider $\mathcal{P} = \{q \in \Delta, q(E) = \alpha\}$ for $\alpha \in (0, 1)$. Then $q^* = \arg\min_{q \in \mathcal{P}} D_f(q\|p)$ if and only if $\frac{dq^*}{dp} = \frac{\alpha}{\theta}\mathbf{1}_E + \frac{1-\alpha}{1-\theta}\mathbf{1}_{E^c}$, where $\theta = p(E)$. Hence $D_f(q^*\|p) = f\left(\frac{\alpha}{\theta}\right)\theta + f\left(\frac{1-\alpha}{1-\theta}\right)(1 - \theta)$.

*Proof.* We adapt an argument similar to that used in Csiszar (1975) for relative entropy. Let $p(E) = \theta$, and consider $q \ll p$ such that $q(E) = \alpha$. Let $r = dq/dp$. The constraint $q(E) = \alpha$ implies that:

$$\int_E r\, dp = \alpha, \qquad \int_{E^c} r\, dp = 1 - \alpha.$$

By Jensen's inequality on $E$ and the convexity of $f$,

$$\int_E f(r)\, dp \geqslant p(E) f\left(\frac{1}{p(E)} \int_E r\, dp\right) = \theta f\left(\frac{\alpha}{\theta}\right).$$

Similarly,

$$\int_{E^c} f(r)\, dp \geqslant (1-\theta) f\left(\frac{1-\alpha}{1-\theta}\right).$$

Therefore

$$D_f(q \| p) \geqslant \theta f\left(\frac{\alpha}{\theta}\right) + (1-\theta) f\left(\frac{1-\alpha}{1-\theta}\right).$$

Strict convexity gives equality iff $r$ is constant on $E$ and on $E^c$, hence iff

$$\frac{dq^*}{dp} = \frac{\alpha}{\theta}\mathbf{1}_E + \frac{1-\alpha}{1-\theta}\mathbf{1}_{E^c}.$$

Thus the closest measure is the Jeffrey update (Jeffrey, 1972): for every event $A$,

$$q^*(A) = \alpha p(A \mid E) + (1-\alpha) p(A \mid E^c)$$

so

$$D(q^* \| p) = \theta f\left(\frac{\alpha}{\theta}\right) + (1-\theta) f\left(\frac{1-\alpha}{1-\theta}\right).$$

□

Plugging in the values in the example, for $L$, the closest element on the boundary of $\mathcal{L}_\epsilon$ to the centre $p^L$ is $q^*$ with $q^*(L) = 0.9 - \epsilon$, so

$$D_f(q^* \| p^L) = 0.9\, f\left(\frac{0.9-\epsilon}{0.9}\right) + 0.1\, f\left(\frac{0.1+\epsilon}{0.1}\right).$$

Similarly, the element of $\mathcal{R}_\epsilon$ to $p^L$ is $q^{**}$ with $q^{**}(R) = 0.1 + \epsilon$, so

$$D_f(q^{**} \| p^L) = 0.1\, f\left(\frac{0.1+\epsilon}{0.1}\right) + 0.9\, f\left(\frac{0.9-\epsilon}{0.9}\right).$$

so $D_f(q^*\|p^L) = D_f(q^{**}\|p^L)$ It follows that the result holds for every $f$-divergence with strictly convex $f$, and hence for the classical statistical distances in Table 4.

□

*Proof of Proposition 2.* Take $o = w_L^L\epsilon^2 + w_B^L(\max\{\epsilon - 0.81, 0\})^2$. $q$, defined by $q(\omega_{LR}) = p^L(\omega_{LR}) - \max\{\epsilon - 0.81, 0\} = 0.09 - \max\{\epsilon - 0.81, 0\}$, $q(\omega_R) = p^L(\omega_R) + \max\{\epsilon - 0.81, 0\} = 0.01 + \max\{\epsilon - 0.81, 0\}$, $q(\omega_L) = p^L(\omega_L) - \min\{\epsilon, 0.81\} = 0.81 - \min\{\epsilon, 0.81\}$ and $q(\omega_N) = p^L(\omega_N) + \min\{\epsilon, 0.81\} = 0.09 + \min\{\epsilon, 0.81\}$, is a probability measure over $\Omega$. Moreover, $\sum_{j=\{L,R,B\}} w_j^L(q(j) - p^L(j))^2 = \left(w_L^L\epsilon^2 + w_B^L(\max\{\epsilon - 0.81, 0\})^2\right)$, so $q \in c^L(o)$. Since, for any $q'$ with $q'(L) < 0.9 - \epsilon$, $\sum_{j=\{L,R,B\}} w_j^L(q'(j) - p^L(j))^2 > \left(w_L^L\epsilon^2 + w_B^L(\max\{\epsilon - 0.81, 0\})^2\right)$, such $q' \notin c^L(o)$, so $c^L(o) \subseteq \mathcal{L}_\epsilon$. For any $\delta \in [0, 0.9]$, consider $q_\delta$ defined by $q_\delta(\omega_{LR}) = p^L(\omega_{LR}) + \max\{0, \delta - 0.09\} = 0.09 + \max\{0, \delta - 0.09\}$, $q(\omega_R) = p^L(\omega_R) + \min\{\delta, 0.09\} = 0.01 + \min\{\delta, 0.09\}$, $q_\delta(\omega_L) = p^L(\omega_L) - \max\{0, \delta - 0.09\} = 0.81 - \max\{0, \delta - 0.09\}$ and $q_\delta(\omega_N) = p^L(\omega_N) - \min\{0.09, \delta\} = 0.09 - \min\{0.09, \delta\}$; this is clearly a probability measure. $\sum_{j=\{L,R,B\}} w_j^L(q_\delta(j) - p^L(j))^2 = \left(w_R^L\delta^2 + w_B^L(\max\{0, \delta - 0.09\})^2\right)$. Noting that $w_R^L\epsilon^2 + w_B^L(\max\{0, \epsilon - 0.09\})^2 < w_L^L\epsilon^2 + w_B^L(\max\{\epsilon - 0.81, 0\})^2$ if and only if $w_B^L\frac{1}{\epsilon^2}\left((\max\{0, \epsilon - 0.09\})^2 - (\max\{\epsilon - 0.81, 0\})^2\right) < w_L^L - w_R^L$, it is straightforward to check that this is the case for all $\epsilon \in [0, 0.9]$ whenever $0.8w_B^L = w_B^L\frac{0.81^2 - 0.09^2}{0.9^2} < w_L^L - w_R^L$. It follows that there exists $\delta > \epsilon$ with $\sum_{j=\{L,R,B\}} w_j^L(q_\delta(j) - p^L(j))^2 \leqslant \left(w_L^L\epsilon^2 + w_B^L(\max\{\epsilon - 0.81, 0\})^2\right) = o$, so $c^L(o) \nsubseteq \mathcal{R}_\epsilon$, as required.

□

Proposition 3 is an immediate corollary of the characterisation of confidence aggregation in Eq. (2) (Proposition D.1) and the observation that the confidence ranking defined in Eq. (3) can equivalently expressed by the implausibility function in Eq. (4). The proof of Proposition 4 is provided in Appendix A.

*Proof of Proposition 5.* The clause of Proposition 5 concerning convex distances is proved precisely as described in Section 3.5. (Note that, since $\Delta$ is compact, and each projection is linear and thus continuous, $P_{\mathcal{P}_1,\ldots\mathcal{P}_m}$ is compact.) It thus remains to prove the part concerning reverse relative entropy.

Under reverse relative entropy, the objective to be minimised in (7) is:

$$\sum_{i=1}^{n}\sum_{j=1}^{m} w_j^i R(p^i|_{\mathcal{P}_j} \| p|_{\mathcal{P}_j}).$$

For each issue $j$, letting $W_j = \sum_{i=1}^{n} w_j^i$ and $\bar{p}_j = \sum_{i=1}^{n} \frac{w_j^i}{W_j} p^i|_{P_j}$, by the calculation in the proof of Proposition 1 (third column),

$$\sum_{i=1}^{n} w_j^i R(p^i|_{P_j} \| p|_{\mathcal{P}_j}) = W_j R(\bar{p}_j \| p|_{\mathcal{P}_j}) + \sum_{i=1}^{n} w_j^i R(p^i|_{P_j} \| \bar{p}_j)$$

where the second term is independent of $p|_{\mathcal{P}_j}$. Hence minimising the objective in (7) is equivalent to minimising

$$\sum_{j=1}^{m} W_j R(\bar{p}_j \| p|_{\mathcal{P}_j})$$

Let $q, r \in \Delta$, and fix $\lambda \in (0,1)$. Since marginalisation is an affine map,

$$(\lambda q + (1-\lambda) r)|_{P_j} = \lambda q|_{P_j} + (1-\lambda) r|_{P_j}.$$

By well-known convexity properties of relative entropy, for all $q, r$ such that $\bar{p}_j$ is absolutely continuous with respect to $q$ and $r$,

$$R\big(\bar{p}_j \big\| \lambda q|_{P_j} + (1-\lambda) r|_{P_j}\big) \leqslant \lambda R(\bar{p}_j \| q|_{P_j}) + (1-\lambda) R(\bar{p}_j \| r|_{P_j}),$$

with equality if and only if

$$q|_{\mathrm{supp}(\bar{p}_j)} = r|_{\mathrm{supp}(\bar{p}_j)}.$$

Summing over $j$ gives

$$\sum_{j=1}^{m} W_j R\big(\bar{p}_j \big\| \lambda q|_{P_j} + (1-\lambda) r|_{P_j}\big) \leqslant \lambda \sum_{j=1}^{m} W_j R\big(\bar{p}_j \big\| q|_{P_j}\big) + (1-\lambda) \sum_{j=1}^{m} W_j R\big(\bar{p}_j \big\| r|_{P_j}\big)$$

and equality holds if and only if

$$q|_{\bigcup_j \mathrm{supp}(\bar{p}_j)} = r|_{\bigcup_j \mathrm{supp}(\bar{p}_j)}$$

By the richness of the partitions (and the fact that $\mathrm{supp}(\bar{p}_j) = \bigcup_{i:w^i_j>0} \mathrm{supp}\, p^i|_{\mathcal{P}_j}$), it follows that there is a unique minimum for (7) on the set

$$S = \left\{ p \in \Delta : \mathrm{supp}\, p \subseteq \bigcup_j \bigcup_{i:w^i_j>0} \mathrm{supp}\, p^i|_{\mathcal{P}_j} \right\}.$$

It thus suffices to show that any optimum will belong to this set. Clearly, if $S = \Delta$, this is the case; so suppose not, and consider $q \in \Delta \backslash S$. If $\bigcap_j \bigcup_{i:w^i_j>0} \mathrm{supp}\, p^i|_{\mathcal{P}_j} \neq \emptyset$, then denote it by $E$. If not, there exists $E = \bigcap_{j\in J} \bigcup_{i:w^i_j>0} \mathrm{supp}\, p^i|_{\mathcal{P}_j}$ for some non-empty $J \subset \{1, \ldots, m\}$ such that $E \cap \bigcup_{j\notin J} \bigcup_{i:w^i_j>0} \mathrm{supp}\, p^i|_{\mathcal{P}_j} = \emptyset$. Let $\alpha = q(S^c) > 0$, $r'$ be the conditionalisation of $q$ on $S^c$ and $r$ be any probability measure with $\mathrm{supp}\, r \subseteq E$. Consider $q' = q + \alpha(r - r')$. By the definition of $r'$, $q' \geqslant 0$, so, since $q, r, r' \in \Delta$, $q' \in \Delta$. Moreover, since, for every $x \in \mathrm{supp}\, q \backslash S$, $q(x) = \alpha r'(x)$, $q' \in S$. For every $j$ such that $E \subseteq \mathrm{supp}\, \bar{p}_j$, $q'(x) \geqslant q(x)$ for all $x \in E$ with strict inequality on a set of non-zero measure, so

$$\begin{aligned} R(\bar{p}_j \| q'|_{\mathcal{P}_j}) &= \int_{\mathcal{P}_j} \log\left(\frac{d\bar{p}_j}{dq'|_{\mathcal{P}_j}}\right) d\bar{p}_j \\ &= \int_{\mathrm{supp}\, \bar{p}_j} \log\left(\frac{d\bar{p}_j}{dq'|_{\mathcal{P}_j}}\right) d\bar{p}_j \\ &< \int_{\mathrm{supp}\, \bar{p}_j} \log\left(\frac{d\bar{p}_j}{dq|_{\mathcal{P}_j}}\right) d\bar{p}_j = R(\bar{p}_j \| q|_{\mathcal{P}_j}) \end{aligned}$$

Moreover, if $E \cap \mathrm{supp}\, \bar{p}_j = \emptyset$, then by a similar argument (noting that, since $q'$ and $q$ differ on an event that receives zero weight under $\bar{p}_j$), $R(\bar{p}_j \| q'|_{\mathcal{P}_j}) = R(\bar{p}_j \| q|_{\mathcal{P}_j})$. Hence

$$\sum_{j=1}^{m} W_j R(\bar{p}_j \| q|_{\mathcal{P}_j}) > \sum_{j=1}^{m} W_j R(\bar{p}_j \| q'|_{\mathcal{P}_j})$$

so the optimum is in $S$, establishing the result.

□

## D.3 Proofs of results in Section 5 and Appendix C

*Proof of Proposition 6.* The proof essentially follows that of Hill (2016, Proposition 3); given the differences in setup (e.g. using strict rather than weak preference as primitive), we include a full proof for completeness.

The implication from the representation to the axioms is straightforward. Consider the implication from the axioms to the representation, and let $(c^{rep}, u, D)$ represent $\succ^{rep}$. For each $f \in \mathcal{A}$, let $\underline{x_f} \in \mathcal{A}^c$ be any $\succ^{rep}$-minimal element of the set of all $y \in \mathcal{A}^c$ such that, for all $x \in \mathcal{A}^c$, if $f \succ^{rep} x$, then $y \succ^{rep} x$. Consider any $f \in \mathcal{A}$. Noting, since $D(x) \leqslant D(f)$ for all $x \in \mathcal{A}^c$ and $f \in \mathcal{A}$, that $\max\{D(f), D(x)\} = D(f)$, it follows from the representation of $\succ^{rep}$ that $u(\underline{x_f}) = \min_{p \in c^{rep}(D(f))} \mathbb{E}_p u(f)$. By representation (9), for any $y \in \mathcal{A}^c$ with $y \prec^{rep} \underline{x_f}$, $f \succ^{rep} y$, whence $f \succ^{dm} y$ by C-Consistency. By the representation, $f \not\succ^{rep} \underline{x_f}$, whence $\underline{x_f} \succeq^{dm} f$ by Caution. Hence, by the continuity of $\succeq^{dm}$, $f \sim^{dm} \underline{x_f}$. Since this holds for all $f \in \mathcal{A}$, $\succeq^{dm}$ is represented by $V(f) = u(\underline{x_f}) = \min_{p \in c^{rep}(D(f))} \mathbb{E}_p u(f)$ as required. □

*Proof of Proposition C.1.* Let $\succeq$ be a variational or unbounded continuous variational preference represented by the variational functional $V$ with $u$ and $\gamma$ (Eq. (C.1)). Let $\Delta$ denote the space of finitely additive probability measures if $\succeq$ is a variational, and the space of countably additive probability measures if it is unbounded continuous. For each $f \in \mathcal{A} \backslash \mathcal{A}^c$, let

$$\mathcal{P}_f = \{p \in \Delta : V(f) = \mathbb{E}_p u(f) + \gamma(p)\}$$

i.e. the set of probability measures achieving the minimum in the variational representation for $f$. For each $f$, by monotonicity, $V(f)$ is finite, so $\mathcal{P}_f \neq \varnothing$, so there exists $m$ such that $\{p \in \Delta : \gamma(p) \leqslant m\} \cap \mathcal{P}_f \neq \varnothing$. If $\succeq$ is variational, then by the Banach-Alaoglu theorem, the space of finitely additive probabilities $\Delta$ is (weakly) compact; since $\{p \in \Delta : \gamma(p) \leqslant m\}$ is weak*-closed, it thus follows that it is (weakly) compact. Similarly, if $\succeq$ is unbounded continuous variational, then by Maccheroni et al. (2006, Theorem 13) , $\{p \in \Delta : \gamma(p) \leqslant m\}$ is (weakly) compact. By the lower semicontinuity of $\gamma$, it follows that it attains its minimum over $\{p \in \Delta : \gamma(p) \leqslant m\} \cap \mathcal{P}_f$, so $\{\gamma(p) : p \in \mathcal{P}_f\}$ has a minimum element, which we denote by $\gamma_f = \min\{\gamma(p) : p \in \mathcal{P}_f\}$.

As a point of notation, let $B(\Omega)$ be the space of real-valued $\Sigma$-measurable simple functions on $\Omega$ under the sup norm and $ba(\Omega)$ be the dual space of bounded and finitely additive set functions, under the total variation norm. $\geqslant$ is the standard ordering on $B(\Omega)$, and $\mathbf{1}$ ($\mathbf{0}$) is the constant function taking value 1 (respectively, 0). Without loss of generality, we assume that $u$ is such that $\mathbf{1}, \mathbf{0} \in u(\mathcal{A}^c)$. Each act $f \in \mathcal{A}$ generates a function $u \circ f \in B(\Omega)$, as standard; let $B(\Omega)^{\mathcal{A}}$ be the image of $\mathcal{A}$ under $u$, and $B(\Omega)^{\mathcal{A}^c}$ be the image of $\mathcal{A}^c$ For any $a \in B(\Omega)^{\mathcal{A}}$, we use $u^{-1}(a) \in \mathcal{A}$ as shorthand for an act mapped to $a$ under $u$.[30]

Let

$$\mathcal{F}_f = \{g \in \mathcal{A}\backslash\mathcal{A}^c : \gamma_g \geqslant \gamma_f\} \cup \left\{ g \in \mathcal{A}\backslash\mathcal{A}^c : \begin{array}{c} \forall \alpha > 0 \text{ s.t. } \alpha(u \circ g) \in B(\Omega)^{\mathcal{A}}, \\ \gamma_{u^{-1}(\alpha(u \circ g))} \leqslant \gamma_f \end{array} \right\}.$$

Consider the set $\mathcal{C}_f \subseteq \Delta$ defined by:

$$\mathcal{C}_f = \bigcap_{g \in \mathcal{F}_f} \{p \in \Delta : \mathbb{E}_p u(g) \geqslant V(g)\}$$

Since it is an intersection of half-spaces, this is convex, closed set. We now establish several properties of it.

**Claim D.9.** For each $f \in \mathcal{A}\backslash\mathcal{A}^c$, $\mathcal{C}_f \neq \varnothing$.

*Proof.* Since, as argued above, for any $m > 0$, $\{p \in \Delta : \gamma(p) \leqslant m\}$ is weakly compact, by the lower semicontinuity of $\gamma$, it attains its minimum. Since $\gamma$ is grounded, it follows that $\{p \in \Delta : \gamma(p) = 0\} \neq \varnothing$. For each $f \in \mathcal{A}\backslash\mathcal{A}^c$, by the variational representation, for each $g \in \mathcal{F}_f$, $\mathbb{E}_p u(g) \geqslant V(g)$ for each $p \in \Delta$ with $\gamma(p) = 0$. Hence $\mathcal{C}_f \supseteq \{p \in \Delta : \gamma(p) = 0\} \neq \varnothing$, so $\mathcal{C}_f$ is non-empty. □

**Claim D.10.** For each $f, g \in \mathcal{A}\backslash\mathcal{A}^c$, if $\gamma_f \geqslant \gamma_g$, then $\mathcal{C}_f \supseteq \mathcal{C}_g$.

*Proof.* Immediate from the definition of $\mathcal{C}_f$. □

[30] All statements in this proof using this notation will hold for all members of the relevant inverse image.

The following fact about variational preferences will prove useful.

**Claim D.11.** For any $a \in B(\Omega)^{\mathcal{A}} \backslash B(\Omega)^{\mathcal{A}^c}$ and $\alpha > 0$ for which $\alpha a \in B(\Omega)^{\mathcal{A}}$, if $\gamma_{\alpha a} \leqslant \gamma_a$, then $\alpha \leqslant 1$. Moreover, if $\alpha \leqslant 1$, $V(u^{-1}(\alpha a)) \geqslant \alpha V(u^{-1}(a))$.

*Proof.* Consider minimal $q, q'$ with $V(u^{-1}(a)) = \mathbb{E}_q a + \gamma(q)$ and $V(u^{-1}(\alpha a)) = \mathbb{E}_{q'} \alpha a + \gamma(q')$. If $\gamma_{\alpha a} = \gamma_a$, the result is immediate, so suppose that $\gamma_{\alpha a} < \gamma_a$. Since $\gamma(q') < \gamma(q)$, $\mathbb{E}_{q'} a > \mathbb{E}_q a$; if not, $V(u^{-1}(a))$ would not be obtained with $q$.

Since $q$ obtains the minimum for $u^{-1}(a)$,

$$\mathbb{E}_q a + \gamma(q) < \mathbb{E}_{q'} a + \gamma(q')$$

whence

$$\gamma(q) - \gamma(q') < \mathbb{E}_{q'} a - \mathbb{E}_q a$$

Similarly, since $q'$ yields the minimum for $u^{-1}(\alpha a)$

$$\gamma(q) - \gamma(q') > \alpha(\mathbb{E}_{q'} a - \mathbb{E}_q a)$$

These two inequalities imply that $\alpha < 1$, as required.

For the second part, for any $\alpha \leqslant 1$:

$$\begin{aligned} V(u^{-1}(\alpha a)) &= \min_{p \in \Delta} \mathbb{E}_p \alpha a + \gamma(p) \\ &\geqslant \min_{p \in \Delta} \mathbb{E}_p \alpha a + \alpha \gamma(p) \\ &= \alpha \left( \min_{p \in \Delta} \mathbb{E}_p a + \gamma(p) \right) = \alpha V(u^{-1}(a)) \end{aligned}$$

establishing the result. □

**Claim D.12.** For each $f \in \mathcal{A} \backslash \mathcal{A}^c$, $\min_{p \in \mathcal{C}_f} \mathbb{E}_p u(f) = V(f)$.

*Proof.* For each $f \in \mathcal{A} \backslash \mathcal{A}^c$, and each $a \in B(\Omega)^{\mathcal{A}} \backslash B(\Omega)^{\mathcal{A}^c}$, let $s^f(a) = \lambda^f(a) a$ where $\lambda^f(a) = \inf\{\alpha > 0 : \gamma_{u^{-1}(\alpha a)} \geqslant \gamma_f\}$ when this set is non-empty and $\lambda^f(a) = 1$ otherwise. Note that, by definition, $s^f(u \circ f) = u \circ f$. Define the functional $I_f : B(\Omega)^{\mathcal{A}} \backslash B(\Omega)^{\mathcal{A}^c} \to \mathbb{R}$ by $I_f(a) = \min_{\substack{p \in \Delta \\ \gamma(p) \leqslant \gamma_f}} \mathbb{E}_p a$. As a maxmin functional, $I_f$ is monotonic, constant additive, positively homogeneous and superadditive. Moreover, we have the following properties.

**Claim D.13.** For every $f \in \mathcal{A} \backslash \mathcal{A}^c$, $I_f(u \circ f) = V(f) - \gamma_f$, and $I_f(s^f(u \circ g)) \geqslant V(u^{-1}(s^f(u \circ g))) - \gamma_f$ for all $g \in \mathcal{A} \backslash \mathcal{A}^c$.

*Proof.* By the definition of $\gamma_f$, there exists $p \in \Delta$ with $\gamma(p) = \gamma_f$ such that $V(f) = \mathbb{E}_p u(f) + \gamma_f$; by the variational representation, for all other $p \in \Delta$ with $\gamma(p) \leqslant \gamma_f$, $V(f) \leqslant \mathbb{E}_p u(f) + \gamma(p) \leqslant \mathbb{E}_p u(f) + \gamma_f$. This establishes the first part of the claim. For the second part, consider $g \in \mathcal{A} \backslash \mathcal{A}^c$ and let $a = s^f(u \circ g)$. We reason by cases.

If $\gamma_{u^{-1}(a)} \geqslant \gamma_f$, then, for every $p \in \Delta$ with $\gamma(p) \leqslant \gamma_f$, $\mathbb{E}_p a + \gamma_f \geqslant \mathbb{E}_p a + \gamma(p) \geqslant V(u^{-1}(a))$, whence $I_f(a) + \gamma_f \geqslant V(u^{-1}(a))$, as required,

If $u^{-1}(a) \in \overline{\{g \in \mathcal{A} \backslash \mathcal{A}^c : \gamma_g \geqslant \gamma_f\}} \backslash \{g \in \mathcal{A} \backslash \mathcal{A}^c : \gamma_g \geqslant \gamma_f\}$, then, for some $p \in \Delta$ with $\gamma(p) < \gamma_f$, $\mathbb{E}_p a + \gamma(p) = V(u^{-1}(a))$, but for every $p \in \Delta$ with $\gamma(p) \leqslant \gamma_f$, $\mathbb{E}_p \alpha a + \gamma_f \geqslant \mathbb{E}_p \alpha a + \gamma(p) \geqslant V(u^{-1}(\alpha a))$, whence $I_f(a) + \gamma_f \geqslant V(u^{-1}(a))$, as required,

Finally, if neither of the previous cases holds, then there is no $\alpha > 1$ with $\gamma_{u^{-1}(\alpha a)} \geqslant \gamma_f$. Hence $V(u^{-1}(a)) = \min_{p \in \Delta : \gamma(p) \leqslant \gamma_f} \mathbb{E}_p a + \gamma(p) \leqslant I_f(a) + \gamma_f$, as required. ☐

We reason using Fan's Theorem (Fan, 1956, Theorem 13; see also the proof of Chateauneuf, 1991, Theorem 2). It implies that, for a family $\{a_j\}_{j \in J}$ of elements in $B(\Omega)$ and corresponding real numbers $\{\alpha_j\}_{j \in J}$, the system:

$$\phi(a_j) \geqslant \alpha_j \qquad \forall j \in J \tag{D.8}$$

has a solution in $ba(\Omega)$ if and only if

$$\sup \sum_{j=1}^{n} \lambda_j \alpha_{k_j} < +\infty \tag{D.9}$$

when $n = 1, 2, 3, \ldots$, $k_j$, $\lambda_j$ vary under the conditions $\lambda_j > 0$ for all $1 \leqslant j \leqslant n$ and $\| \sum_{j=1}^{n} \lambda_j a_{k_j} \| = 1$.

To establish the claim, it suffices to show that, for $f \in \mathcal{A}\backslash\mathcal{A}^c$, with $a_f = u \circ f$, there exists $\phi \in ba(\Omega)$ such that:

$$\begin{aligned} \phi(a_f) \geqslant V(f), &\qquad \phi(-a_f) \geqslant -V(f), \\ \phi(\mathbf{1}) \geqslant 1, &\qquad \phi(-\mathbf{1}) \geqslant -1, \\ \phi(u \circ g) \geqslant V(g) &\quad \forall g \in \mathcal{F}_f, \\ \phi(a) \geqslant 0 &\quad \forall a \in B(\Omega)^{\mathcal{A}}, a \geqslant \mathbf{0} \end{aligned} \tag{D.10}$$

Note first, by Claim D.9, $\mathcal{C}_f$ is non-empty, so there is a solution for the set of inequalities (D.10'), obtained from (D.10) by removing the inequality $\phi(-a_f) \geqslant -V(f)$.

Clearly, the supremum in (D.9) for (D.10) is equal to the maximum of the supremum across all vectors $(\lambda_1, \ldots \lambda_n)$, $(a_f, -a_f, \mathbf{1}, -\mathbf{1}, u \circ g_{k_5}, \ldots u \circ g_{k_m}, a_{m+1}, \ldots, a_n)$ satisfying the conditions stated after (D.9) such that $\lambda_1 - \lambda_2 \geqslant 0$, and the supremum across all vectors of this set such that $\lambda_1 - \lambda_2 < 0$. It suffices thus to show that each of these suprema is finite.

Consider a pair of vectors $(\lambda_1, \ldots \lambda_n)$, $(a_f, -a_f, \mathbf{1}, -\mathbf{1}, u \circ g_{k_5}, \ldots u \circ g_{k_m}, a_{m+1}, \ldots, a_n)$ involved in the supremum (D.9) for the system (D.10), i.e. such that $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \geqslant 0$, $\lambda_j > 0$ for all $5 \leqslant \lambda \leqslant n$, and

$$\left\| \lambda_1 a_f + \lambda_2(-a_f) + \lambda_3 \mathbf{1} + \lambda_4(-\mathbf{1}) + \sum_{j=5}^{m} \lambda_j (u \circ g_{k_j}) + \sum_{k=m+1}^{n} \lambda_k a_k \right\| = 1,$$

There is clearly a one-to-one map between the subset of such vectors with $\lambda_1 - \lambda_2 \geqslant 0$ and pairs $(\lambda_1 - \lambda_2, \lambda_3 \ldots \lambda_n)$, $(a_f, \mathbf{1}, -\mathbf{1}, u \circ g_{k_5}, \ldots u \circ g_{k_m}, a_{m+1}, \ldots, a_n)$ involved in the supremum (D.9) for (D.10'); moreover, the corresponding expressions in (D.9) are identical. Since, as noted above, the set of inequalities (D.10') has a solution, by Fan's Theorem this supremum is finite, so the supremum (D.9) for the system (D.10) over all vectors with $\lambda_1 - \lambda_2 \geqslant 0$ is finite.

Now consider pairs of vectors $(\lambda_1, \ldots \lambda_n)$, $(a_f, -a_f, \mathbf{1}, -\mathbf{1}, u \circ g_{k_5}, \ldots u \circ g_{k_m}, a_{m+1}, \ldots, a_n)$ involved in the supremum (D.9) for the system (D.10), with $\lambda_1 - \lambda_2 < 0$. It follows from

$$\left\| \lambda_1 a_f + \lambda_2(-a_f) + \lambda_3 \mathbf{1} + \lambda_4(-\mathbf{1}) + \sum_{j=5}^{m} \lambda_j (u \circ g_{k_j}) + \sum_{k=m+1}^{n} \lambda_k a_k \right\| = 1,$$

that

$$(\lambda_1 - \lambda_2) a_f + (\lambda_3 - \lambda_4)\mathbf{1} + \sum_{j=5}^{m} \lambda_j (u \circ g_{k_j}) + \sum_{k=m+1}^{n} \lambda_k a_k \leqslant \mathbf{1}$$

This holds if and only if:

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1}\mathbf{1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1}(u \circ g_{k_j}) + \sum_{k=m+1}^{n} \frac{\lambda_k}{\lambda_2 - \lambda_1} a_k \leqslant a_f$$

Since $a_k \geqslant \mathbf{0}$, this implies that

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1}\mathbf{1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1}(u \circ g_{k_j}) \leqslant a_f \tag{D.11}$$

By the monotonicity, constant additivity and convexity of the variational representation $(V \circ u)$, if $\frac{\lambda_j}{\lambda_2 - \lambda_1} < 1$ for all $j = 5, \ldots, m$ and $\sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \leqslant 1$, then it follows that:

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} V(g_{k_j}) \leqslant V(f)$$

Now consider the case where this does not hold; in this case, $\sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} > 1$, since all $\lambda_j > 0$. Letting $a_{k_j} = u \circ g_{k_j}$ for each $j$, (D.11) implies that

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1}\mathbf{1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} s^f(a_{k_j}) \leqslant a_f$$

Since each $g_{k_j} \in \mathcal{F}_f$, it follows from the definition of $s^f$ and Claim D.11 that $\frac{1}{\lambda^f(a_{k_j})} \geqslant 1$ for each $j$. By the monotonicity, constant additivity, positive homogeneity and superadditivity of $I_f$, the previous relation implies that:

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} I_f(s^f(a_{k_j})) \leqslant I_f(a_f)$$

By Claim D.13, this implies that

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} \left(V(u^{-1}(s^f(a_{k_j})) - \gamma_f\right) \leqslant V(f) - \gamma_f$$

whence

$$\begin{aligned} &\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} V(u^{-1}(s^f(a_{k_j})) \\ < &\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} V(u^{-1}(s^f(a_{k_j})) + \gamma_f \left(1 - \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})}\right) \\ \leqslant & V(f) \end{aligned}$$

since, by the previous observations, $\sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} \frac{1}{\lambda^f(a_{k_j})} > 1$. Since $s^f(a_{k_j}) = \lambda^f(a_{k_j}) a_{k_j}$ for all $j$, where $\lambda^f(a_{k_j}) \leqslant 1$, it follows from Claim D.11 that

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} V(u^{-1}(a_{k_j})) \leqslant V(f)$$

whence:

$$\frac{\lambda_3 - \lambda_4 - 1}{\lambda_2 - \lambda_1} + \sum_{j=5}^{m} \frac{\lambda_j}{\lambda_2 - \lambda_1} V(g_{k_j}) \leqslant V(f)$$

So

$$\lambda_1 V(f) + \lambda_2(-V(f)) + \lambda_3 - \lambda_4 + \sum_{j=5}^{m} \lambda_j V(g_{k_j}) \leqslant 1,$$

and hence the supremum (D.9) for (D.10) over all vectors with $\lambda_1 - \lambda_2 < 0$ is finite.

Hence, the supremum in (D.9) for (D.10) is finite. Thus, by Fan's Theorem, the set of equations (D.10) has a solution, i.e. there exists $p \in \mathcal{C}_f$

with $\mathbb{E}_p u(f) = V(f)$, as required. $\square$

**Claim D.14.** For any decreasing sequence $\gamma_i \to \gamma$, $\bigcap_i \bigcap_{f:\gamma_f \geqslant \gamma_i} \mathcal{C}_f = \bigcap_{f:\gamma_f \geqslant \gamma} \mathcal{C}_f$.

*Proof.* For any decreasing sequence $\gamma_i \to \gamma$,

$$\bigcap_i \bigcap_{f:\gamma_f \geqslant \gamma_i} \mathcal{C}_f = \bigcap_{f:\gamma_f > \gamma} \mathcal{C}_f$$

Since each of the $\mathcal{C}_f$ are closed and nested, this is a closed set; clearly, $\bigcap_{f:\gamma_f \geqslant \gamma} \mathcal{C}_f$ is too. It thus suffices to show that $\bigcap_{f:\gamma_f > \gamma} \mathcal{C}_f = \bigcap_{f:\gamma_f \geqslant \gamma} \mathcal{C}_f$. Suppose not for reductio, and let $p \in \bigcap_{f:\gamma_f > \gamma} \mathcal{C}_f \backslash \bigcap_{f:\gamma_f \geqslant \gamma} \mathcal{C}_f$. By a standard separating hyperplane argument, there exists $a \in B(\Omega)^{\mathcal{A}} \backslash B(\Omega)^{\mathcal{A}^c}$ and $x, \epsilon \in \mathbb{R}_{>0}$ with $\mathbb{E}_p a \leqslant x - \epsilon$ and $\mathbb{E}_q a \geqslant x$ for all $q \in \bigcap_{f:\gamma_f \geqslant \gamma} \mathcal{C}_f$. It follows from the latter inequalities and Claim D.12 that $\gamma_{u^{-1}(a)} \leqslant \gamma$. By the definition of $\mathcal{C}_f$, if $\gamma_{u^{-1}(\alpha a)} \leqslant \gamma$ for all $\alpha > 0$, then $\mathbb{E}_q a \geqslant x$ for all $q \in \mathcal{C}_f$ for every $f$ with $\gamma_f > \gamma$, contradicting the existence of $p$. So there exists $\alpha' > 1$ with $\gamma_{u^{-1}(\alpha' a)} > \gamma$. Pick any $b = \alpha a$ with $\gamma_{u^{-1}(b)} > \gamma$ and consider $\{\beta b : \beta \in (0, 1)\}$. By the previous observation, there exists a $\beta b$ in this set with $\gamma_{u^{-1}(\beta b)} \leqslant \gamma$, let $b' = \beta' b$ be the supremum of these elements. By Claims D.12 and D.10 and the specification of $p$, $a$ and $x$, for any increasing sequence $\beta_i \to \beta'$, for each $\beta_i$, $V(u^{-1}(\beta_i b)) = \min_{q \in \mathcal{C}_{u^{-1}(\beta_i b)}} \mathbb{E}_q \beta_i b \geqslant u^{-1}(\beta_i \alpha x)$; by the continuity of $V$, it follows that $V(u^{-1}(\beta' b)) \geqslant \beta' \alpha x$. Moreover, $\mathbb{E}_p \beta' b \leqslant \beta' \alpha x - \beta' \alpha \epsilon$. Hence, for any decreasing sequence $\beta_i \to \beta'$, for each $\beta_i$ by the fact that $\beta'$ is a supremum, $\gamma_{u^{-1}(\beta_i b)} > \gamma$, so, by Claim D.12, $V(u^{-1}(\beta_i b)) = \min_{q \in \mathcal{C}_{u^{-1}(\beta_i b)}} \mathbb{E}_q \beta_i b \leqslant \beta_i \alpha x - \beta_i \alpha \epsilon$, whence, by the continuity of $V$, $V(u^{-1}(\beta' b)) \leqslant \beta' \alpha x - \beta' \alpha \epsilon$ contradicting the previous inequality. Hence there exists no such $p$, and the result is established. $\square$

Let $O = [0, \sup_{p \in \Delta} \gamma(p))$, define $c(o) = \bigcap_{f:\gamma_f \geqslant o} \mathcal{C}_f$. By the definition and the previously noted facts, $c$ is a convex, closed confidence ranking. Define $D : \mathcal{A} \to O$ by $D(f) = \gamma_f$. It remains to show that $(c, D, u)$ represents $\succeq$ according to (11). This is immediate for constant acts, so consider $f \in \mathcal{A} \backslash \mathcal{A}^c$. By the definitions of $c$ and $D$, and Claims D.10 and

D.12:

$$\begin{aligned}\min_{p\in c(D(f))} \mathbb{E}_p u(f) &= \min_{\bigcap_{g\in\mathcal{F}_f}\{p\in\Delta:\mathbb{E}_p u(g)\geqslant V(g)\}} \mathbb{E}_p u(f)\\ &= V(f)\end{aligned}$$

as required. □

*Proof of Proposition C.2.* For the first clause, by the axiomatisation in Maccheroni et al. (2006), it clearly suffices to show that $\succeq^{dm}$ satisfies Weak C-Independence and Uncertainty Aversion (see Maccheroni et al., 2006, Section 3 for definitions).

**Weak Certainty Independence** Suppose that $f_\alpha x \succeq^{dm} g_\alpha x$, for $f, g \in \mathcal{A}$, $x, y \in \mathcal{A}^c$ and $\alpha \in (0,1]$. Then, by the representation:

$$\min_{c(D(f_\alpha x))} \mathbb{E}_p u(f) \geqslant \min_{c(D(g_\alpha x))} \mathbb{E}_p u(g)$$

If $f, g \notin \mathcal{A}^c$, then by Assumption C.1, $D(f_\alpha y) = D(f_\alpha x)$ and $D(g_\alpha y) = D(g_\alpha x)$ whence the previous inequality holds for $D(f_\alpha y)$ and $D(g_\alpha y)$, so $f_\alpha y \succeq^{dm} g_\alpha y$, as required. Similar arguments establish the other cases.

**Uncertainty Aversion** Suppose that $f \sim^{dm} g \sim^{dm} x$ for $x \in \mathcal{A}^c$. By the representation, for every $y \in \mathcal{A}^c$ with $y \prec^{dm} x$, $f \succ^{dm} y$ and $g \succ^{dm} y$, so by Assumption C.2, $D(f_\alpha g)$ is such that $\min_{p\in c(D(f_\alpha g))} \mathbb{E}_p u(f_\alpha g) > u(y)$, whence $f_\alpha g \succ^{dm} y$. It follows from continuity of $\succ^{dm}$ that $f_\alpha g \succeq^{dm} x$, as required.

For converse implication, suppose that $c$ is strictly increasing.

**Assumption C.1** Consider $f \in \mathcal{A}\backslash\mathcal{A}^c$, $x, y \in \mathcal{A}^c$ and $\alpha \in (0,1]$, and let $z \in \mathcal{A}^c$ be such that $f_\alpha x \sim^{dm} z_\alpha x$ (such $z$ exists by standard arguments). Since $\succeq^{dm}$ satisfies Weak C-Independence, $f_\alpha y \sim^{dm} z_\alpha y$. Plugging this into representation (11), it follows that:

$$\min_{c(D(f_\alpha x))} \mathbb{E}_p u(f) = \min_{c(D(f_\alpha y))} \mathbb{E}_p u(f)$$

which, given that $c$ is strongly increasing and $f \notin \mathcal{A}^c$, can only hold if $D(f_\alpha x) = D(f_\alpha y)$ as required.

**Assumption C.2** Consider $f, g \in \mathcal{A}$, $\alpha \in [0,1]$, $x \in \mathcal{A}^c$ and $o \in O$, and suppose that $f \succ^{dm} x$, $g \succ^{dm} x$ and $\min_{p \in c(o)} \mathbb{E}_p u(f_\alpha g) \leqslant u(x)$. Since variational preferences are concave, it follows that $f_\alpha g \succ^{dm} x$, whence, by the representation (11), $\min_{p \in c(D(f_\alpha g))} \mathbb{E}_p u(f_\alpha g) > u(x)$. Since $c$ is increasing, it follows that $o > D(f_\alpha g)$, as required. $\square$