Confidence, consensus and aggregation*

Brian Hill
CNRS & HEC Paris†

April 26, 2024

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

Abstract

This paper develops and defends a new approach to belief aggregation, involving confidence in beliefs. It is axiomatically characterised by a variant of the Pareto condition that enjoins respecting consensuses borne of compromise. Confidence aggregation generalises standard probability aggregation rules—such as linear pooling—whilst avoiding the spurious unanimity issues that have plagued them. It generates the first family of probability aggregation rules that can faithfully accommodate within-person expertise diversity, hence resolving a longstanding challenge. It is dynamically rational, insofar as it commutes with update. Finally, it recovers as special cases both Bayesian and non-Bayesian approaches to model misspecification.

Keywords: Belief aggregation, confidence in beliefs, Pareto principle, linear pooling, spurious unanimity, expertise, consensus, model averaging, model misspecification.

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 a collection of honest and well-intentioned experts' beliefs be aggregated into a set of group judgements? Doubtless the most popular proposal in economics, statistics and risk analysis is linear pooling, which takes a weighted average of probabilistic beliefs (e.g. Stone, 1961; Cooke, 1991; Gilboa et al., 2004). It is based on a principle of consensus preservation: any consensus in beliefs concerning a particular issue, or in preferences depending on that issue, is preserved in the group beliefs or preferences. This *issue-wise* consensus preservation is formulated by the Pareto principle underpinning some preference-based axiomatisations of linear pooling (Mongin, 1995), as well as of generalisations to non-Bayesian decision models (Crès et al., 2011; Danan et al., 2016). However, it has recently faced severe challenges.

One central problem comes in examples where there is unfounded consensus on an issue, or *spurious unanimity* (Mongin, 2016). In such cases, linear pooling respects the issue-wide consensus, despite its spuriousness. For instance, consider 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 both 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 linear pool of their judgements is given in the final row of the table. Irrespective of the weights assigned to the individuals, it preserves their common judgement on the effect on both sectors—a consequence of the Pareto principle in this context. 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 a limited effect on real estate; Ray does so because of the low probability he

	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 has a limited effect on the sector(s) in the top row

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

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 *issue-wide consensuses* is unjustified (Mongin, 2016), and hence a problem for linear pooling (Bradley, 2017b; Mongin and Pivato, 2020; Dietrich, 2021). The stated aim of respecting consensus is clearly reasonable; the problem, it seems, is that linear pooling sometimes respects the wrong consensuses.

The example also illustrates a second, apparently distinct challenge, involving the way linear pooling, as well as popular alternatives including geometric pooling, incorporates expertise. It does 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*. It thus cannot reflect expertise differences across issues: for instance, it 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 pooling rules in the literature, does not allow this.

Both challenges are significant for the committee's decision in this example. If it follows linear pooling and accepts the 'spurious' consensus that the probability of a limited effect on both sectors is low, it would not implement the rise. By contrast, if it considered each expert's judgement on their respective sectors of expertise, this would suggest a higher probability of a limited effect on both, hence allowing for the possibility of the rise. Moreover, the decision-relevant factor—whether there is a limited effect on both sectors—lies at the intersection of the committee members' fields of expertise, hence posing the problem of how to incorporate their different levels of expertise across issues.

This paper proposes a new approach to belief aggregation that incorporates respect for consensus into rationally-founded aggregation—hence retaining the gist of the Pareto principle—whilst avoiding commitment to unfounded or spurious consensuses. As a byproduct and separate contribution, the approach naturally accommodates within-person cross-issue expertise diversity.

Our approach introduces two novel insights. For the first, note that spuriousness arises in examples where issue-level consensus is respected to the detriment of other elements of agents' states of opinion, such as information, other beliefs, reasons or evidence (Mongin and Pivato, 2020; Dietrich, 2021; Bommier et al., 2021). Presumably Laura's and Ray's similar judgements on the 'Both' issue are based on different evidence, supporting the low probabilities they assign to Real Estate and Labour respectively. It thus seems that an agent's declared probability for an event does not exhaust her relevant judgements pertaining to that event. This echoes a position defended in the literatures on belief representation and decision under uncertainty: a probability judgement does not fully capture all relevant aspects of a (rational) agent's state of belief concerning an event. For instance, several approaches (e.g. Marinacci, 2015; Maccheroni et al., 2006; Chateauneuf and Faro, 2009; Hill, 2013, 2019b; Bradley, 2017a) model belief states as comprising agents' *confidence in beliefs*. To the extent that one's confidence in a belief is related to one's evidence, information and reasons underlying it (Hill, 2019a), confidence could serve as an overarching concept to refer to what is being overlooked by linear pooling in these spuriousness examples.

Our second insight concerns consensus: if issue-wise consensus preservation is problematic, what sort of consensus should be preserved instead? We recognise that consensus typically requires *compromise*. One often speaks of achieving a consensus, through which agents may compromise on or 'put aside' some opinions to retain others. Under this conception, a consensus is not a single issue on which people happen to have the same beliefs, but a common ground comprising of a coherent set or 'corpus' of positions acceptable to all. More precisely, such a *corpus-level* consensus is a coherent set of judgements, each emanating from some member of the group, and such that each member would be ready to 'set aside'—or compromise—any potential disagreements in the interests of the consensus. Note that a corpus may be more or less complete: the associated judgements need not settle every question. It seems reasonable that the judgements holding in such corpus-level consensuses be preserved in the group's beliefs.

But what compromises would agents be willing to make to achieve consensus? In reply, our approach weaves together the two previous insights by invoking confidence as a determinant of the propensity to compromise. A (rational) individual is surely more concerned in seeing a judgement held with high confidence respected in the final group beliefs, even if that is at the expense of some lower-confidence beliefs. This suggests that confidence determines compromise via the following maxim: the more confident an individual is in a belief, the less willing she is to compromise on that belief.

The first contribution of this paper is to propose an aggregation rule for confidence in beliefs that preserves corpus-level consensus judgements, where consensuses are borne of compromise regulated by confidence according to this maxim. We provide preference-based axiomatic foundations for the rule, showing that it is characterised by a Pareto-style axiom, which essentially states that judgements in such consensuses are preserved.

Our second main set of contributions concerns the aforementioned challenges to standard pooling rules. We first show that popular probabilistic opinion pooling rules can be reproduced as special cases of our *confidence aggregation* rule, corresponding to particular assumptions on individuals' confidence in their beliefs. This sheds light on the comparison with existing approaches: whereas classic pooling rules essentially postulate what individuals are willing to compromise to arrive at group beliefs, our approach uses precisely the compromises determined by the individuals themselves, as encoded in the confidence they have in their beliefs.

This analysis sets the stage for the integration of within-person cross-issue expertise diversity. An individual with more expertise on one issue than another would be justified in having more confidence *ceteris paribus* in her beliefs concerning the former issue. Drawing on this insight, we explore the consequences of our aggregation rule when applied in cases involving different degrees of confidence—reflecting differing expertise—according to the issue under consideration. It yields group judgements that more strongly respect an individual's judgement on the issues on which she is an expert, and less so on those on which she has less expertise. Beyond establishing that our approach resolves the expertise challenge, these examples show that it does not respect spurious issue-level consensuses resulting from ignoring expertise differences. Hence it resolves the spurious unanimity challenge too.

As an application, we use confidence aggregation to generate a new family of probabilistic belief aggregation rules that can accommodate within-person expertise diversity. To our knowledge, these are the first such rules in the literature, and certainly the first to have received preference-theoretic axiomatic foundations.

Moreover, we briefly consider the issue of dynamic rationality, which is typically evoked to justify geometric pooling (Genest and Zidek, 1986; Dietrich, 2021). Drawing on a recently proposed account of rational update for confidence in belief (Hill, 2022), we show that confidence aggregation fully satisfies dynamically rationality with respect to this update, in the standard sense: the two commute.

Finally, we consider the application of our approach to the analogous question of deciding in the face of multiple 'expert' models. It recovers as special cases both the Bayesian standard—Bayesian Model Averaging—and the main non-expected utility approaches to model misspecification, whilst suggesting improvements to incorporate within-model expertise diversity.

Section 2 sets out the framework and the aggregation rule. Section 3, which contains examples illustrating the approach, shows how confidence aggregation overcomes the challenges to linear pooling and generates a new family of probability aggregation rules tailored to cases of within-person expertise diversity. Section 4 contains a preference-based characterisation of confidence aggregation, and Section 5 considers its dynamic rationality. Section 6 applies it to decisions with models, and concern for misspecification. Section 7 discusses remaining related literature. Proofs and supplementary material are contained in the Appendices.

2 Confidence aggregation

2.1 Preliminaries

Setup Let Ω be a non-empty set of states.[1] Subsets of Ω are called *events*; *partitions* are sets of mutually disjoint events whose union is Ω . For any partition $\mathcal{P}$ (including Ω itself), $\Delta(\mathcal{P})$ denotes the set of probability measures over $\mathcal{P}$; henceforth, we let $\Delta = \Delta(\Omega)$.[2] For any $p \in \Delta$ and partition $\mathcal{P}_j$, $p|_{\mathcal{P}_j} \in \Delta(\mathcal{P}_j)$ denotes the projection of p into $\Delta(\mathcal{P}_j)$.

A *(statistical) distance* ρ on $\Delta(\mathcal{P})$ is a function $\rho : \Delta(\mathcal{P})^2 \to [0, \infty]$ such that: $\rho(q, p) = 0$ if and only if $p = q$; and $\rho(\bullet, q)$ is a lower semicontinuous function, for all $q \in \Delta(\mathcal{P})$. A distance ρ is *convex* if, for every $q \in \Delta(\mathcal{P})$, the function $\rho(\bullet, q)$ is strictly convex.[3] A (convex) *classical statistical distance* d is the specification, for each partition $\mathcal{P}$ (including Ω 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 $\Delta(\mathcal{P})$. Metrics, such as the Euclidean metric (for finite Ω), and divergences, such as the relative entropy or Kullback-Leibler divergence, are (convex) classical statistical distances, insofar as they specify a distance for each probability space (Table 2).

$O \subseteq \mathbb{R}$ 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 a closed left-bounded interval in $\mathbb{R}$, with the associated topology.[4] We shall 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 profiles: $\mathbf{o} \geqslant \mathbf{o}'$ 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 ambiguity. The belief state of an agent—incorporating confidence—

[1] For the purposes of exposition, Ω can be taken to be finite, though extension to the infinite case is straightforward.

[2] Throughout, we take the weak* topology on Δ and $\Delta(\mathcal{P})$.

[3] That is, for all $p, r \in \Delta$ with $p \neq r$ and $\alpha \in (0, 1)$, $\rho(\alpha p + (1-\alpha)r, q) < \alpha\rho(p, q) + (1-\alpha)\rho(r, q)$.

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

is represented by a *confidence ranking*: a function $c : O \to 2^{\Delta}\backslash\varnothing$ that is increasing in the containment order on sets and is upper semicontinuous.[5] For each confidence level o, $c(o)$ is the set of priors representing the beliefs the agent hold with confidence of at least o. For any $o \in O$ and increasing, upper semicontinuous function $c : \{o' \in O : o' \geqslant o\} \to 2^{\Delta}\backslash\varnothing$, the *natural extension* of c, denoted $\bar{c}$, is the confidence ranking defined by $\bar{c}(o') = c(o')$ for $o' \geqslant o$ and $\bar{c}(o') = c(o)$ otherwise.

The *centre* of confidence ranking c is its smallest element, i.e. $\min_{o \in O} c(o)$. A confidence ranking c is *centred* if its centre is a singleton. Centred confidence rankings represent Bayesians with confidence: agents who assign a precise probability to every event (namely, that given by the centre), though may have more confidence in some judgements than others (as represented by the rest of the confidence ranking). A confidence ranking c is *convex* (respectively, *closed*) if, for every $o \in O$, $c(o)$ is a convex (resp. closed) set. For a confidence ranking c, its *convex closure* c^{clconv} 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.

Confidence rankings admit two alternative equivalent representations. Firstly, note that each probability judgement—i.e. judgement such as 'the probability of A is greater than x', 'A is probabilistically independent of B' etc.—corresponds to a subset of Δ, namely the set of probability measures where the judgement holds. Noting this, the function $conf : 2^{\Delta} \to O \cup \{\varnothing\}$, defined by:

$$conf(\mathcal{J}) = \begin{cases} \varnothing & \min_{o \in O} c(o) \nsubseteq \mathcal{J} \\ \max\{o : c(o) \subseteq \mathcal{J}\} & \text{otherwise} \end{cases} \tag{1}$$

picks out, for any probability judgement $\mathcal{J}$, the agent's confidence in $\mathcal{J}$—the largest confidence level at which $\mathcal{J}$ is held if it is held, and nothing otherwise. A confidence ranking also generates a unique *implausibility function* $\iota : \Delta \to O \cup \varnothing$ defined by $\iota(p) = \min\{o \in O : p \in c(o)\}$ whenever the set is non-empty, and $\iota(p) = \varnothing$ otherwise.[6] This yields the 'implausibility' of each probability measure, in terms of the smallest confidence level such

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

[6] This is well-defined by the upper semicontinuity of c.

that the probability measure doesn't contradict a judgement held with that much confidence.[7]

We consider a group of n individuals, indexed by i; individual 0 is the group. A tuple $(c^1, \ldots, c^n)$ of confidence rankings for each individual, where c^i is the confidence ranking of individual i, is called a *profile*. The group confidence ranking is denoted c^0. As noted, this can equivalently be written as a profile of implausibility functions $(\iota^1, \ldots, \iota^n)$ and group implausibility function ι^0.

Related models and distance-based confidence As just noted, the representation of confidence used here is equivalent to a real-valued function on the space of probability measures. As such, it includes many prominent models of decisions under uncertainty—such as smooth, variational, multiplier and confidence preferences (Klibanoff et al., 2005; Maccheroni et al., 2006; Hansen and Sargent, 2001; Chateauneuf and Faro, 2009)—which involve such functions (or functions generating them) in their preference representation, and often interpret them in terms of confidence. It also includes the weaker representation of confidence in beliefs developed by Hill (2013, 2019b) under the calibration in Hill (2019a) (see Section 7).

Moreover, the alternative representation in terms of implausibility implies that it is possible to generate a confidence ranking from a probability measure and a (statistical) distance on Δ. More specifically, given a probability $p \in \Delta$, weight $w \in \mathbb{R}_{>0}$ and distance ρ on Δ, the following defines the confidence ranking centred on p with associated implausibility function $\iota(q) = w\rho(q, p)$.

Definition 1. Let $p \in \Delta$ be a probability measure, ρ a statistical distance on Δ and $w \in \mathbb{R}_{>0}$. The *w confidence ranking generated by p under ρ*—or simply the *w ρ-confidence ranking generated by p*—is defined by $c(o) = \{q \in \Delta : wd(q, p) \leqslant o\}$ for all $o \in O$.

As an illustration, Table 2 lists some well-known classical statistical distances, and the corresponding generated confidence rankings. Several of the previous references use such distance-generated confidence rankings; for in-

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

Generating distance	$\rho(q, p) =$	w ρ-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	$R(p\|q)$	$\{q \in \Delta : wR(p\|q) \leqslant o\}$

Table 2: Examples of convex classical distances and distance-generated confidence rankings.
Note: Euclidean distance only well-defined on finite Ω. R is the relative entropy, defined by: $R(p\|q) = -\sum p(\omega)(\log \frac{q(\omega)}{p(\omega)})$.

stance, the w relative entropy confidence ranking is involved in multiplier preferences (Hansen and Sargent, 2001; Maccheroni et al., 2006).

2.2 Consensus-preserving confidence aggregation

To introduce the notion of consensus, consider a tuple $\mathbf{o} = (o_1, \ldots, o_n)$ of confidence levels and a profile $(c^1, \ldots, c^n)$ of confidence rankings. If $\bigcap_i c^i(o_i) = \varnothing$, then the individuals' beliefs at the confidence levels $\mathbf{o}$ are in contradiction. By contrast, 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 other words, when $\bigcap_i c^i(o_i) \neq \varnothing$, it represents a corpus-level consensus, in which a probability judgement holds if and only if it is held by at least one individual at the confidence level specified for them by $\mathbf{o}$.[8]

In the consensus characterised by $\bigcap_i c^i(o_i)$, individuals are not compromising on the beliefs they hold with confidence $\mathbf{o}$ or more: these are all retained. Rather, each individual i compromises by only putting her beliefs held with confidence o_i or more 'on the table', and ignoring any lower-confidence beliefs. To that extent, the compromises involved in such a consensus are reflected in the confidence level each individual uses to

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

determine the beliefs they contribute. When higher confidence levels are involved, more compromise is required by the individuals. However, this means that the resulting consensus is more robust: it only contains judgements on which individuals are particularly confident.

There may be several such consensuses, involving different compromises—different levels of confidence required in particular individuals' beliefs for them to be taken into account. To translate them into levels of confidence deemed relevant for the group, we use a *confidence-level aggregator*: an operator $\otimes : O^n \to O$ that is monotonic in each argument, i.e. such that for every pair of profiles of confidence levels with $\mathbf{o} \geqslant \mathbf{o}'$, $\otimes\mathbf{o} \geqslant \otimes\mathbf{o}'$. For a consensus obtained with individual confidence levels $\mathbf{o}$, the confidence-level aggregator picks out the group confidence warranted in the associated consensus judgements. Monotonicity reflects the fact that the higher the individual confidence levels $\mathbf{o}$ behind the consensus, the higher the corresponding group confidence level. Since higher individual confidence levels translate into a consensus involving more compromise, but that is also more robust, this seems reasonable.

In our preference-based characterisation (Section 4), the relevant confidence-level aggregator will be endogenous; however, it may be instructive to consider some examples.

Example 2.1 (Affine aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \sum_{i=1}^{n} w_i o_i + \chi$ for $w_i \in \mathbb{R}_{>0}$, $\chi \in \mathbb{R}$ is called an **affine aggregator**.

Example 2.2 (Average aggregator)**.** The special case of the affine aggregators with the same weights are **average aggregators**: $\otimes\mathbf{o} = \sum \frac{1}{n} o_i + \chi$ for χ as above.

Example 2.3 (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, is called a **maximum aggregator**.

Example 2.4 (Minimum aggregator)**.** An aggregator of the form $\otimes\mathbf{o} = \psi(\min\{o_i\})$, with $\psi : O \to O$ as above, is called a **minimum aggregator**.

We can now introduce our aggregation rule. Since individuals' beliefs are represented by confidence rankings, a suitable aggregation rule needs to relate the profile of individual confidence rankings with a group confidence ranking. Each confidence-level aggregator $\otimes$ generates such a rule, in the form of the function $F_\otimes$, defined as follows. For every profile $(c'^1, \ldots, c'^n)$ of confidence rankings, $F_\otimes(c'^1, \ldots, c'^n) = \overline{\Phi_\otimes(c'^1, \ldots, c'^n)}$, where, for every $o \in O$ such that $\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o} \bigcap_i c'^i(o_i) \neq \varnothing$

$$\Phi_\otimes(c'^1, \ldots, c'^n)(o) = \bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o} \bigcap_{i=1}^{n} c'^i(o_i) \tag{2}$$

For the purposes of the preference-based characterisation in Section 4, where we follow the economic literature and work in a single-profile setup (e.g. Mongin, 1995; Gilboa et al., 2004; Crès et al., 2011; Danan et al., 2016), this yields the following definition of consensus-preserving confidence aggregation for a fixed confidence ranking c^0 and profile $(c^1, \ldots, c^n)$.

Definition 2. The group confidence ranking c^0 is a *consensus-preserving confidence aggregation of* $(c^1, \ldots, c^n)$ if there exists a confidence-level aggregator $\otimes$ such that $c^0 = F_\otimes(c^1, \ldots, c^n)$. In this case, we say that c^0 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 group forms judgements with confidence level o by looking at the consensuses considered to warrant a confidence level o or less according to $\otimes$.[9] More specifically, the group holds a probability judgement with confidence o if that judgement holds for all such consensuses: this is guaranteed by the union in Eq. (2). In that sense, it preserves those judgements that hold unanimously across the appropriate consensuses. In the resulting group beliefs, none of the judgements held at confidence level o contradict the corresponding consensus judgements, though if two consensuses contradict each other on an issue, neither's judgement will be retained in the group beliefs with confidence o.

A noteworthy consequence of this aggregation rule is that group and individual confidence in a judgement co-vary: because of the monotonicity of $\otimes$, the group confidence in a judgement is higher when the individual

[9]The use of consensuses corresponding to confidence levels less than and equal to o in Eq. (2) ensures that c^0 is a well-defined confidence ranking, without requiring any assumptions on $\otimes$. As discussed in Appendix B, it can be replaced by the union over consensuses with confidence level *equal to o* for various notable families of $\otimes$, including those in the examples above.

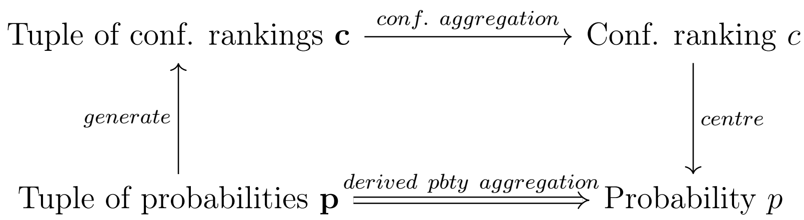

Figure 1: Using confidence aggregation to generate probability aggregation rules

beliefs feeding into the relevant consensuses are held at higher confidence levels. This is arguably a reasonable property for a procedure for aggregating beliefs and confidence.

As shown in Proposition C.1 (Appendix C.1), this aggregation rule can be formulated in terms of implausibility functions: c^0 is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$ under $\otimes$ if and only if, for all $p \in \Delta$

$$\iota^0(p) = \begin{cases} \otimes(\iota^1(p), \ldots, \iota^n(p)) & \text{if } \forall i \ \iota^i(p) \in O \\ \varnothing & \text{otherwise.} \end{cases} \tag{3}$$

3 Confidence, probability aggregation and expertise diversity

We now consider how confidence aggregation deals with the challenges to linear pooling. We first show that standard probability aggregation rules can be recovered as special cases of confidence aggregation; analysing the underlying assumptions will naturally reveal how confidence aggregation can resolve both the spurious unanimity and the within-person cross-issue expertise diversity challenges. The discussion also contains several examples of confidence aggregation, and culminates in a new probability aggregation rule for within-person expertise diversity.

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 distance and a weight are specified, each probability measure generates a unique centred confidence ranking (Definition 1, Section 2.1). This provides the following possibility for using consensus-preserving confidence aggregation to aggregate probability measures. Given a profile of probability measures, take a profile of confidence rankings generated by them, say 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 c is centred, then this is in fact a singleton, and the procedure yields a unique probability measure. 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 n-tuple of weights, with $w^i \geqslant 0$ for all i, with strict inequality for some i. For each row in Table 3, the following holds:*

(*) Let c be the consensus-preserving confidence aggregation under an average confidence-level aggregator of w^i confidence rankings generated by p_i under the distance given in the first column of Table 3. Then its centre is 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 satisfies the equation in the third column of the Table.

Hence the two most prominent pooling rules in the literature (Genest and Zidek, 1986; Mongin, 1995; Dietrich, 2021) in fact correspond to special cases of confidence aggregation, where the probability measures involved in the rules are the centres of the individuals' and group's confidence rankings. 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 confidence rankings for the individuals

Generating distance	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	$p = \sum_i \frac{w^i}{\sum_{i=1}^n w^i} p_i$

Table 3: 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.

in the group. As is clear from the comparison of the cases in Table 3, the 'shape' of the confidence ranking determines the pooling rule reproduced. In this sense, the use of, say, linear pooling, can be thought of as amounting to the assumption that individuals' confidence rankings are generated by the Euclidean or reverse relative entropy distances.[10] And the evaluation of this pooling rule can thus pass via an appraisal of the corresponding assumption.

The result also suggests a new strategy for facing the challenges cited in the Introduction. If the weaknesses of traditional pooling rules can be connected to how the confidence rankings are generated, then applying confidence aggregation with different generation methods may avoid them.

3.2 Representing expertise using confidence rankings

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.

[10]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 the individuals' belief states—and, as shall be clear below, evaluate them as such.

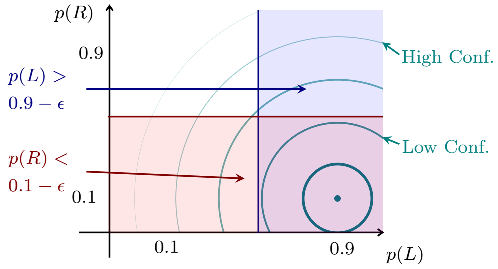

(a) Illustration of 'issue-neutrality'.
Note: The blue area represents the probability judgement, $\mathcal{L}_\epsilon$, that $p(L)$ is within ϵ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement, $\mathcal{R}_\epsilon$, that $p(R)$ is within ϵ of $p^L(R) = 0.1$.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

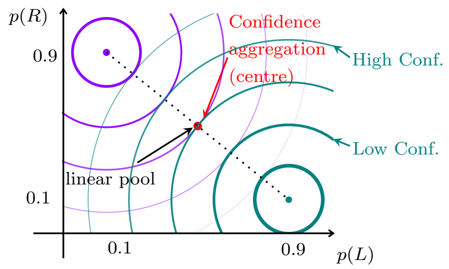

(b) Illustration of Proposition 1.
Note: The red point is the centre of the result of confidence aggregation applied to the two confidence rankings (Proposition 1). Each point on the dotted line is obtained by linear pooling (with some choice of weights). This graph displays the case of $w^L = w^R$;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

Example 3.1. We formalise the example from the Introduction with a state space $\Omega = \{\omega_{LR}, \omega_L, \omega_R, \omega_N\}$ where ω_{LR} (respectively ω_L, ω_R, ω_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 $p^L(\omega_{LR}) = 0.09$, $p^L(\omega_L) = 0.81$, $p^L(\omega_R) = 0.01, p^L(\omega_N) = 0.09$. So, 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 $0.1 + \epsilon$.[11] 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 C.3 (Appendix C.2) shows that, under the two confidence-ranking-generating procedures yielding linear pooling (Table 2), the confidence in the two judgements is the same, no matter the ϵ. Figure 2a illustrates the intuition: given the 'circular' shape 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, but not on the issue concerned by the judgement—labour or real estate.

The confidence rankings generating standard pooling rules thus represent individuals as having the same confidence in the probability judgements encoded in their probability measure p^i, no matter the issues that these judgements concern. As such, they cannot properly capture an individual who has 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 distances on the probability space, assume that there is no 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 group be-

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

liefs. We now confirm this suggestion, and show how it can produce new expertise-sensitive probability aggregation rules.

For the presentation, we focus on issues that can be related to events in Ω; see Section 3.5 for a generalisation. Consider a sequence $\mathcal{P}_1, \ldots \mathcal{P}_m$ of partitions of Ω; 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 $\mathcal{P}_1, \ldots \mathcal{P}_m$ is *rich* if, for any $(p_1, \ldots, p_m) \in \prod_{j=1}^{m} \Delta(\mathcal{P}_j)$, there exists at most one $p \in \Delta$ with $p|_{\mathcal{P}_j} = p_j$ for all $j = 1, \ldots, m$. 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 Ω, so this sequence of partitions is rich.

Now consider the following family of centred confidence rankings.

Definition 3. Let $\mathcal{P}_1, \ldots \mathcal{P}_m$ be partitions of Ω 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 $o \in O$,

$$c(o) = \left\{ q \in \Delta : \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \leqslant o \right\} \tag{4}$$

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 1, with implausibility function determined by the following distance on Δ:

$$\iota(q) = \rho_d^{\mathbf{w}}(q,p) = \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \tag{5}$$

Apart from the special case involving a single partition $\mathcal{P} = \Omega$, $\rho_d^{\mathbf{w}}$ are not classical distances.[12]

The issue-specific weights in $\mathbf{w}$ d-confidence rankings can capture an agent's 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 suppose that Laura's confidence ranking is generated by it with Euclidean distance and vector of weights $\mathbf{w^L} = (w_L^L, w_R^L, w_B^L)$. I.e. Laura has the confidence ranking:

$$c^L(o) = \left\{ q \in \Delta : \sum_{j=\{L,R,B\}} 2w_j^L (q(j) - p^L(j))^2 \leqslant o \right\} \tag{6}$$

The weights reflect Laura's 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*, she is represented as having more confidence in judgements concerning issues with higher weights. This is borne out by the following proposition.

Proposition 2. *Suppose, in (6), 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 ϵ below its best-guess probability 0.9 is held with more confidence than a judgement about R that involves the same divergence ϵ 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.

[12] In particular, $\rho_d^{\mathbf{w}}$ does not automatically provide a distance on any space except Ω.

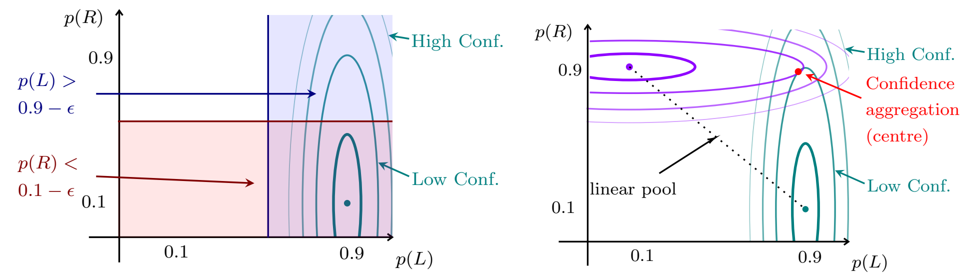

(a) Illustration of Proposition 2.
Note: As in Figure 2a, the blue area represents the probability judgement, $\mathcal{L}_\epsilon$, that $p(L)$ is within ϵ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement, $\mathcal{R}_\epsilon$, that $p(R)$ is within ϵ of $p^L(R) = 0.1$. 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 4). 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. (6).
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. (6)—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.

So $w^L_L > w^L_R$ 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^L_B 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 shortly.

If Laura can be naturally represented by a confidence ranking of the form (6), Ray can be represented with a similar confidence ranking, centred on p^R as specified in Table 1,[13] with weights $\mathbf{w^R} = (w^R_L, w^R_R, w^R_B)$ where

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

$w^R_R > w^R_L$, translating his relative expertise in real estate.

3.3 Confidence aggregation with within-person expertise diversity

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

Proposition 3. *Suppose that each agent $i = 1, \ldots, n$ has a confidence ranking of the form (4), with classical distance d, centre p^i and vector of positive real-valued weights $\mathbf{w^i}$. Then the centre of the consensus-preserving confidence aggregation under an average confidence-level aggregator is:*

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

This result is an immediate corollary of the characterisation of confidence aggregation in Eq. (3) (Proposition C.1) and the observation that the confidence ranking defined in Eq. (4) can equivalently expressed by the implausibility function in Eq. (5).

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

Example 3.4. Suppose that Laura and Ray have 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 (6), 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 $\frac{w^L_B + w^R_B}{w^L_R + w^R_R}$. Given that, in the example, Laura's expertise concerns L but not specifically R or $B = L \cap R$, and similarly for Ray, it seems reasonable that these ratios will be low.

As shown in Appendix A (see also Section 3.4), the aggregate confidence ranking is centred, and, when $\frac{w_B^L + w_B^R}{w_L^L + w_L^R} \to 0$ and $\frac{w_B^L + w_B^R}{w_R^L + w_R^R} \to 0$—i.e. the confidence in judgements concerning B is dwarfed by the overall confidence in the judgements concerning L and R—the centre probabilities tend to:

$$
\begin{aligned}
p(L) &\to \frac{w_L^L}{w_L^L + w_L^R} p^L(L) + \frac{w_L^R}{w_L^L + w_L^R} p^R(L) \\
p(R) &\to \frac{w_R^L}{w_R^L + w_R^R} p^L(R) + \frac{w_R^R}{w_R^L + w_R^R} p^R(R) \\
p(B) &\to \begin{cases} 0.09 & \text{if } p(L) + p(R) - 1 \leqslant 0.09 \\ p(L) + p(R) - 1 & \text{otherwise} \end{cases}
\end{aligned}
$$

So the centre probability for L, $p(L)$, 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. If, as the example suggests, Laura has more expertise than Ray on the labour market, so $w_L^L > w_L^R$, this probability for L will be closer to Laura's ($p^L(L)$), as one would have wanted. Similarly, $p(R)$ tends to the weighted average of the individuals' judgements about R, except that here the weights corresponding to the issue R are involved. Since Ray is more of a specialist here, his weight will be larger $w_R^R > w_R^L$, so the centre judgement will be closer to his judgement on R ($p^R(R)$). 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 cross-issue expertise differences, thus yields group judgements that follow each individual more closely on their area of expertise. As such, it fairs better on this score than linear (or, for that matter, geometric) pooling.

Given that the centres of the confidence rankings generated as in (6) are probability measures with $p(B) = 0.09$, the centre of the aggregate ranking will stick as close to this value as possible. If the weights yield issue-wide weighted averages which are consistent with $p(B) = 0.09$, then this is the value of $p(B)$. If not, as will typically be the case, then $p(B)$ takes the value closest to 0.09 that satisfies the constraints, i.e. $p(L) + p(R) - 1$. Since this is typically not 0.09,[14] this example demonstrates that the confidence aggregation rule does not respect spurious unanimities.

We by no means wish to suggest that all individuals' confidence rankings subscribe to the form in Definition 3. The aim of this example is rather to illustrate a simple way in which the confidence approach can capture within-person expertise diversity, and to show that confidence aggregation faithfully reflects these expertise differences in the resulting group beliefs. As such, it resolves the within-person expertise diversity challenge. Moreover, the example also shows that our aggregation procedures avoid the much-discussed problem with spurious unanimities: when the individuals are not comparatively confident in their judgements about B—so the agreement is indeed spurious—the common judgement is not adopted by the group.

Since the expertise-sensitive generation of confidence rankings in Definition 3 only requires probability measures as input, one could potentially use it in tandem with confidence aggregation to define a new pooling rule. We now consider this possibility.

3.4 Expertise-sensitive pooling

Just as confidence aggregation can recoup standard pooling rules by using confidence rankings generated by classical distances (Proposition 1; Figure 1), we now show that using $\mathbf{w}^i$ d-confidence rankings generates a new family of pooling rules. To this end, let us define the function yielding the centre in the result of confidence aggregation applied to $\mathbf{w^i}$ d-confidence rankings.

Definition 4. Let $\mathcal{P}_1, \ldots, \mathcal{P}_m$ be a set of partitions, and d a classical distance. The function $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m} : \Delta^n \to 2^\Delta$ is defined by

$$F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}(p^1, \ldots, p^n) = \arg\min_{p\in\Delta} \sum_{i=1}^{n} \sum_{j=1}^{m} w^i_j d(p|_{\mathcal{P}_j}, p^i|_{\mathcal{P}_j}) \qquad (8)$$

[14] E.g. when $w^L_L = w^R_R = 0.75$, $w^L_R = w^L_R = 0.25$, $p(L) = p(R) = 0.7$, and $p(B) = 0.4$.

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 measure. However, the optimisation problem defining $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ can typically be reduced to a recognisable 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, each of which is derived from some probability measure on Ω. Note that, since projection is a linear map, $P_{\mathcal{P}_1,\ldots\mathcal{P}_m}$ is a convex set. Moreover, it is typically defined by a collection of inequalities. For instance, in the case of our running example (Example 3.2), $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$ is defined by the following linear inequalities imposed by the fact that $B = L \cap R$: for any $(p_L, p_R, p_B) \in \Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$

$$\begin{aligned} p_L(L) &\geqslant p_B(B) \\ p_R(R) &\geqslant p_B(B) \\ 1 &\geqslant p_L(L) + p_R(R) - p_B(B) \end{aligned} \tag{9}$$

The centre of the aggregate confidence ranking (8) 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:

$$\underset{(p_1,\ldots,p_m)\in P_{\mathcal{P}_1,\ldots,\mathcal{P}_m}}{\arg\min} \sum_{i=1}^{n}\sum_{j=1}^{m} w_j^i d(p_j, p^i|_{\mathcal{P}_j}) \tag{10}$$

For the $\mathbf{w}$ Euclidean-confidence rankings in Example 3.3, this is a quadratic optimisation problem over a convex set (see Appendix A for details). More generally, whenever d is convex, (10) is a minimisation of a strictly convex lower semicontinuous function on a convex set, so there is a unique minimum. So whenever $\mathcal{P}_1, \ldots, \mathcal{P}_m$ is rich, (10) defines a unique probability measure in Δ. Hence, for each convex d and rich set of issues, $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}$ is single-valued. This establishes the following Proposition.

Proposition 4. *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 Δ^n to Δ.*

Confidence aggregation thus generates this new well-defined pooling rule, which we call *expertise-sensitive pooling*. Since this is essentially the rule used in Example 3.4, all of the conclusions there—and in particular the capacity to naturally reflect within-person expertise diversity in the group judgement—apply equally for this pooling rule. Indeed, as discussed in Section 7, the limit expressions concerning L and R in Example 3.4 are reminiscent of early suggestions in the pooling literature; unlike them, however, expertise-sensitive pooling is well-defined.

3.5 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 3.5. 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.[15] This can be reflected using a vector of weights $\mathbf{w^L} = (w^L_L, w^L_R, w^L_I)$ (resp. $\mathbf{w^R} = (w^R_L, w^R_R, w^R_I)$) and the following confidence ranking:

$$c^L_{Ind}(o) = \left\{ q \in \Delta : \begin{array}{c} \sum_{j=\{L,R\}} 2w^L_j(q(j) - p^L(j))^2 \\ +2w^L_I\left(q(B) - q(L).q(R)\right)^2 \end{array} \leqslant o \right\} \tag{11}$$

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'

[15] 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.

from independence are contained in the set of priors, translating the limited confidence in independence.

The solution of the minimisation problem (10) can be obtained similarly to the analysis in Example 3.4, yielding as centre of the aggregate confidence ranking p with:

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

Here the aggregation on each of the issues L and R uses issue-specific weights, reflecting differing confidence, as in the limit case in Example 3.4. For the issue B, agents' beliefs concerning the independence of L and R generates the probability.

In tandem with the preceding discussion, 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.[16] 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 (11) shows how confidence aggregation can respect independence, whilst retaining much of the spirit of linear pooling. For instance, when $w_L^L = w_R^L = w_L^R = w_R^R$, the resulting centre probability is $p^{LP}(L) = 0.5$, $p^{LP}(R) = 0.5$, $p^{LP}(B) = 0.25$: i.e. the same as linear pooling for the issues L and R, but with independence retained (and hence a different B).

[16] Note that whilst these examples used confidence rankings based on the Euclidean distance, similar techniques can be applied to other distances, such as relative entropy.

The beliefs about the independence of L and R in Example 3.5 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. (11), 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.

4 Characterising Confidence Aggregation

In this section we provide a preference-based axiomatisation of consensus-preserving confidence aggregation (Definition 2) in a single-profile setting. First, we set out the decision framework and preference representation.

4.1 Preferences

Preliminaries Consider a standard Anscombe-Aumann (1963)-style framework, as adapted by Fishburn (1970). Let $\mathcal{X}$, the set of *consequences*, be a convex subset of a vector space; for instance it could be the set of lotteries over a set of prizes. $\mathcal{A}$ is the set of *acts*: (measurable) functions from states Ω 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 α-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 use $\succ$ (perhaps with superscripts) to denote a strict preference relation on $\mathcal{A}$. 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 $V(f) > V(g)$.

Each individual and the group has a preference relation $\succ^i$: the tuple $(\succ^1, \ldots, \succ^n)$ is a profile of individual preference relations, and $\succ^0$ is the group preference.

Decision models As noted in Section 2.1, the representation of confidence in beliefs used here is compatible with several models of decision under uncertainty. Here we work with the confidence family, which requires the weakest properties of the belief representation and supports both ambiguity averse and incomplete preference models (Hill, 2019b). Under the ambiguity averse maxmin-EU model in this family (Hill, 2013), preferences are represented by

$$\min_{p \in c(D(f))} \mathbb{E}_p u(f) \tag{12}$$

whereas a typical corresponding incomplete preference model (Hill, 2016) is such that for all acts $f, g \in \mathcal{A}$, $f \succ g$ if and only if:

$$\mathbb{E}_p u(f) > \mathbb{E}_p u(g) \qquad \text{for all } p \in c(\max\{D(f), D(g)\}). \tag{13}$$

In these expressions $\mathbb{E}_p$ is the expectation with respect to a probability measure $p \in \Delta$,[17] $u : \mathcal{X} \to \mathbb{R}$ is a non-constant affine utility function, c is a closed convex confidence ranking and $D : \mathcal{A} \to O$ satisfies the following *richness* condition: for every $f, g \in \mathcal{A}$ and $o \in D(\mathcal{A})$, there exists $h \in \mathcal{A}$ and $\alpha \in (0, 1]$ such that $\max\{D(f_\alpha h), D(g_\alpha h)\} = o$. This function, called the *cautiousness coefficient*, picks out the confidence level the decision maker considers relevant for evaluating each act, and hence each decision. As shown in the cited papers, it captures the decision maker's attitudes to choosing on the basis of limited confidence (or ambiguity attitudes). Models (12) and (13) are related in the standard way (Ghirardato et al., 2004; Gilboa et al., 2010) ; see Hill (2013, 2016, 2019b) for discussion and details.

If (12) or (13) holds, we say that (c, D, u) represents $\succ$ (under the relevant model). In each case, the representing u is unique up to positive affine transformation, $c(O)$ is unique up to convex closure, and $c \circ D$ is unique.

Following Danan et al. (2016), we study aggregation in the context of incomplete preferences, and thus assume that all preferences are represented according the previous incomplete preference model.[18]

[17] I.e. for any $\phi : \Omega \to \mathbb{R}$, $\mathbb{E}_p \phi = \sum_{\omega \in \Omega} p(\omega)\phi(\omega)$ and similarly for infinite Ω.

[18] To the extent that the decision models cited in Section 2.1 involve complete preferences with representations generating confidence rankings, this incomplete preference model can be thought of as embedded in them. Accordingly, our axiomatic analysis of aggregation can be considered as applying under those models (see also Danan et al., 2016). A corresponding characterisation of confidence aggregation can be easily obtained

Assumption 1. *For each $i = 0, \ldots, n$, $\succ^i$ is represented according to (13).*

To focus on aggregation of beliefs, we follow the literature (e.g. Crès et al., 2011) in assuming that all individuals and the group have the same tastes. Since the confidence model has two parameters representing tastes—the utility function u and the cautiousness coefficient D—this is expressed by the following assumption.[19]

Assumption 2. *Let (c^0, D, u) represent $\succ^0$ according to (13). Then, for each $i = 1, \ldots, n$, there exists c^i with (c^i, D, u) representing $\succ^i$ according to (13).*

Henceforth, we fix representations (c^i, D, u) of $\succ^i$, for $i = 0, \ldots, n$.

Stakes A central idea behind the confidence family is that the beliefs one relies on to decide are held to a level of confidence that is appropriate given the importance of the decision (Hill, 2013, 2016, 2019b; Bradley, 2017a). For instance, (13) represents decision makers for which determinate preferences held at low stakes—where less confidence is required—may become indeterminate at higher stakes. In the light of this, when higher-confidence beliefs are invoked—i.e. $\max\{D(f), D(g)\} > \max\{D(f'), D(g')\}$—then this is an indication that the decision maker 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 (13), this can be formalised by a surjective function $\sigma : \mathcal{A} \times \mathcal{A} \to \mathcal{S} \subseteq \mathbb{R}$, assigning to each binary choice the stakes involved in it. Hill (2016) contains several examples of such (real-valued) notions of stakes.[20] For $(f, g) \in \mathcal{A} \times \mathcal{A}$ and $s \in \mathcal{S}$, we say that (f, g) *has stakes s* if $\sigma(f, g) = s$. Assumption 2 guarantees that all individuals and the group

under (12), relying on the aforementioned relationship between the models.

[19] Given the aforementioned 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 and the previous one can be reformulated in behavioural terms, drawing, for instance, on the choice-based foundations for the weak preference version of (13) provided by Hill (2016) (which can be extended to strict preference using the work of Bewley, 1986; Karni, 2011).

[20] 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 (13) (the stakes in the choice as the highest stakes among available acts), Hill (2016) contains examples of others (for which the results below also hold).

have preferences consistent with a single stakes function σ, in the sense that D and σ are linked by a monotone transformation.[21]

Given a preference relation $\succ$ represented according to (13) and a stakes level $s \in \mathcal{S}$, the derived relation $\succ_s$ is defined as follows: for all $f, g \in \mathcal{A}$, $f \succ_s g$ 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 $f_\alpha h \succ g_\alpha h$.[22] 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 on the Democrat candidate winning the 2024 US President election, yielding \$1 million if you win and a loss of \$1 million if not, and nothing g. The other choice is between a similar bet f' on the 2028 election, with stakes (winnings and losses) a million times less in utility terms, and no utility change, g'. An agent with beliefs that are more precise and slightly more favorable for the 2024 bet might nevertheless choose the bet in the 2028 choice but have indeterminate preferences in the 2024 one because of the difference in stakes: with lower stakes, he can rely on low-confidence beliefs when comparing f' and g', but not for the choice between f and g. However, if the 2024 choice was evaluated at the low stakes level, say s, then f would typically be chosen over g: i.e. $f \succ_s g$. 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 We denote tuples of stakes levels $(s_1, \ldots, s_n) \in \mathcal{S}^n$ with vectors $\mathbf{s}$. Under this notation, s_i is understood to be the i^{th} stakes level under $\mathbf{s}$. The following definition shall prove crucial.

Definition 5. For a profile of stakes levels $\mathbf{s} = (s_1, \ldots s_n) \in \mathcal{S}^n$, define the relation $\succ_{\mathbf{s}}$ on $\mathcal{A}$ by $\succ_{\mathbf{s}} = \bigcup_{i=1}^n \succ_{s_i}^i$. $\mathbf{s}$ *exhibits consensus* when $\succ_{\mathbf{s}}$ is not contradictory, and it does not exhibit consensus otherwise. Moreover, we say that $\succ^0$ *respects the consensus* $\succ_{\mathbf{s}}$ *at stakes level* s if $\mathbf{s}$ exhibits consensus and $\succ_s^0 \subseteq \succ_{\mathbf{s}}$.

The relation $\succ_{\mathbf{s}}$ assembles all the (determinate) preferences of the individuals in the group, at the stakes levels specified by $\mathbf{s}$. The group ex-

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

[22]This is well-defined because of the richness of D.

hibits consensus across $\mathbf{s}$ if none of the assembled preferences contradict each other, in the sense of strictly preferring different acts. In other words, were each individual i to only put their preferences at stakes level s^i 'on the table', a coherent consensus position would exist, consisting of all such preferences. In this case, $\succ_{\mathbf{s}}$ represents the preferences under this consensus. The group preference $\succ^0$ respects the consensus $\succ_{\mathbf{s}}$ at a given stakes level s if it doesn't decide more than that consensus: all of the preferences decided upon in $\succ^0_s$ appear in the consensus, though some determinate preferences in the consensus may be left open in $\succ^0_s$. In other words, consensus respect at stakes level s means that the group doesn't adopt stronger positions on preferences than the consensus, at that stakes level.

4.2 Confidence aggregation and Pareto

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

Axiom (Issue-wise Pareto)**.** *For all acts $f, g \in \mathcal{A}$, if $f \succ^i g$ for all i, then $f \succ^0 g$.*

As discussed in the Introduction, this principle encodes respect for issue-wise consensus, and hence faces challenges relating to spurious unanimity. We thus consider the following variant.

Axiom (Corpus-wise Pareto)**.** *For every stakes level $s \in \mathcal{S}$ and acts $f, g \in \mathcal{A}$, if $f \succ_{\mathbf{s}} g$ for all $\mathbf{s}$ for which $\succ^0$ respects the consensus at s, then $f \succ^0_s g$.*

Rather than asking the group to adopt a preference if everyone in the group holds it, Corpus-wise Pareto looks at whether it holds in all relevant consensuses. If the preference holds in all consensuses respected at a given stakes level, then the group adopts that preference at those stakes. 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 Issue-wise Pareto nor Corpus-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 could be considered more general.

Our characterisation requires two auxiliary axioms.

Axiom (Consensus-based beliefs)**.** *For every stakes level $s \in \mathcal{S}$ and acts $f, g \in \mathcal{A}$, if $f \not\succ^0_{s'} g$ for every stakes level s' such that some consensus $\succ_{\mathbf{s}}$ is respected at s', then $f \not\succ^0_s g$.*

Axiom (Non-degeneracy)**.** *There exists a tuple of stakes levels* $\mathbf{s}$ *exhibiting consensus.*

In aggregation, groups beliefs should come from individuals' beliefs. Under confidence aggregation, the latter translate into group beliefs principally in the context of corpus-level consensuses. In terms of preferences, this occurs at stakes levels where some consensus is respected. Consensus-based beliefs states that all group 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, they can come to a consensus.

We have the following characterisation result.

Theorem 1. *Let $\{\succ^i\}, \succ^0$ satisfy Assumptions 1 and 2. They satisfy Corpus-wise Pareto, Consensus-based beliefs and Non-degeneracy if and only if, up to convex closure, c^0 is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$.*

Moreover, there is a unique minimal confidence-level aggregator $\otimes$ under which c^0 is a consensus-preserving confidence aggregation: that is, for all $\otimes'$ such that c^0 is a consensus-preserving confidence aggregation of $(c^1, \ldots, c^n)$ under $\otimes'$, $\otimes'(\mathbf{o}) \geqslant \otimes(\mathbf{o})$ for all $\mathbf{o}$ such that $\bigcap_{i=1}^n c^i(o_i) \neq \varnothing$.

So the central axiom characterising confidence aggregation is Corpus-wise Pareto, which is no more than a reformulation of the standard Pareto condition to apply to (corpus-level) consensuses rather than individual preferences. Indeed, Consensus-based beliefs can be dropped whenever the group confidence ranking is centred. More generally, without it, the group confidence ranking is always that obtained by a confidence aggregation, except at confidence levels at the bottom of the ranking.

No assumption of a particular confidence-level aggregator is required for this result; rather, the appropriate aggregator is determined endogenously by the individual and group preferences. Moreover, there is a unique minimal one: that is, one which always takes the lowest value across all aggregators representing the profile of preferences. Further axioms can be added to characterise the special cases of confidence-level aggregators mentioned in Section 2.2; details are provided in Appendix B.

5 Dynamic rationality

A common theme in the literature is the interaction between aggregation and update. Dietrich (2021) argues that a 'rational group' requires belief aggregation to be in sync with updating. This is typically formulated in terms of commutivity between the two: aggregation followed by update on some information yields the same group beliefs as updating all individual beliefs on the information and then aggregating. The version of this condition for Bayesian beliefs, where updating is performed on events (or likelihoods) by Bayesian conditionalisation, has been called 'external Bayesianism' in the pooling literature (Genest and Zidek, 1986) or 'Dynamic Rationality' by Dietrich (2021).

However, the natural domain for our aggregation approach is not Bayesian beliefs but richer and more refined confidence in beliefs. Here, Bayesian conditionalisation no longer applies, without revision. Hill (2022) proposes a *confidence update* rule for the general representation of confidence in beliefs used here, and argues for its normative validity, suggesting in particular that it deals appropriately with situations where Bayesian update struggles. So the question of dynamic rationality in our context is whether confidence aggregation commutes with confidence update.

In the framework set out in Section 2.1, the probability-threshold confidence update rule from Hill (2022, Definition 2) can be defined as follows, where, for a set $\mathcal{C} \in 2^{\Delta}\backslash\varnothing$ and event E, $\mathcal{C}_E = \{p(\bullet|E) : p \in \mathcal{C}, p(E) > 0\}$, and a probability-threshold function ρ_E is a decreasing function $O \to [0, 1]$.

Definition 6 (Confidence Update)**.** For event $E \subseteq 2^{\Delta}\backslash\varnothing$, confidence ranking $c : O \rightarrow 2^{\Delta}\backslash\varnothing$ and probability-threshold function $\rho_E : O \rightarrow [0,1]$, the confidence update of c by E under ρ_E is the ranking $c|\rho_E = \overline{\Phi}$, where the partial function $\Phi : O \rightarrow 2^{\Delta}\backslash\varnothing$ is defined, for all $o \in O$ such that $\{p \in c(o) : p(E) \geqslant \rho_E(o)\} \neq \varnothing$, by:

$$\Phi(o) = \{p \in c(o) : p(E) \geqslant \rho_E(o)\}_E \tag{14}$$

See Hill (2022) for a full discussion and axiomatic characterisation of this and a more general class of confidence update rules.

We have the following result (where $F_{\otimes}$, the confidence aggregation rule with confidence-level aggregator $\otimes$, is as defined in Section 2.2).

Theorem 2. *For every tuple of confidence rankings* $(c^1, \ldots, c^n)$*, every confidence-level aggregator* $\otimes$*, every event* E *and probability-threshold function for it* ρ_E*:*

$$F_{\otimes}(c^1|\rho_E, \ldots, c^n|\rho_E) = F_{\otimes}(c^1, \ldots, c^n)|\rho_E \tag{15}$$

So confidence aggregation commutes with confidence update: it is 'dynamically rational', to use Dietrich's (2021) term. Such coherence has been argued to be an important property of an aggregation rule, so much so that some use it to promote aggregation rules having this property, and to criticise those that don't. Theorem 2 thus provides a reassuring message concerning confidence aggregation's credentials on this score.

6 Model misspecification and confidence

The focus thus far has been on the contribution that confidence aggregation can make to challenges faced by probability aggregation *per se*. Now we briefly comment on its relevance in the context of subsequent group decision, including in situations involving non-expected utility preferences. A topical relevant such case is where the 'experts' are models.

To this end, consider a decision maker faced with a set $\mathcal{M} \subset \Delta$ of models, each of which, like the experts in standard probability aggregation, is a probability measure over states. As set out in Section 3, confidence aggregation can be applied to confidence rankings generated from $\mathcal{M}$, given an assignment $w : \mathcal{M} \to \Re_{\geqslant 0}$ of weights to models, a distance ρ, and a confidence-level aggregator $\otimes_{m\in\mathcal{M}}$ (Section 2.2). As noted in Section 2.1, the resulting confidence representation is consistent with a range of decision models; for illustration we consider the ambiguity averse model (12) in the confidence family (Section 4; Hill, 2013, 2019b). Applied on the aggregate confidence ranking, it evaluates an act f according to

$$\min_{\substack{q\in\Delta:\\ \otimes_{m\in\mathcal{M}} w(m)\rho(q,m)\leqslant D(f)}} \mathbb{E}_q u(f) \tag{16}$$

where the notation is as in Section 4. Like Gilboa and Schmeidler (1989) maxmin-EU, this decision rule evaluates acts by the worst-case expected utility over a set of priors; unlike it, the set used is determined by the act evaluated, with acts involving higher stakes being evaluated using sets corresponding to higher confidence levels. As such, this rule recovers several approaches in the literature as special cases corresponding to various settings of $\otimes$, ρ and the stakes of the group's decision (Table 4).

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 is a singleton for low enough $D(f)$), then (16) coincides with subjective expected utility (SEU). So, for confidence-level aggregators $\otimes$ and distances ρ yielding the results of, say, linear pooling (Proposition 1), (16) coincides with SEU 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 4, row 1).

Medium / High stakes: Model misspecification At higher stakes, sets further up the confidence ranking will feature in (16). For a fixed $D(f)$, (16) generalises Hansen and Sargent's (2001; 2008) constraint preferences, which are the special case with singleton $\mathcal{M}$ and relative entropy distance ρ.

To that extent, confidence aggregation combined with confidence decision rule (12) naturally reflects concern for model misspecification. Indeed, just as constraint preferences yield the same optima as so-called multiplier preferences on various classes of decision problems (Hansen and Sargent, 2008), for convex distances and confidence-level aggregators, the optimal choice under (16) in these problems coincides with the optimal choice under:[23]

$$\min_{q\in\Delta} \left(\mathbb{E}_q u(f) + \lambda \otimes_{m\in\mathcal{M}} w(m)\rho(q,m)\right)\right) \tag{17}$$

for appropriate λ. (17) is a generalisation of multiplier preferences (which involve singleton $\mathcal{M}$ and relative entropy ρ; Hansen and Sargent, 2001), using the same ingredients as in (16) and providing the same solutions on the aforementioned classes of problems. As such, it dovetails with recent literature on multi-model generalisations of multiplier preferences; as set out in Table 4 (rows 3 & 4), several recent models are special cases of (17), corresponding to particular settings of $\otimes$ and ρ.

Whilst misspecification-motivated models in the literature bake everything together into the decision rule, the confidence-aggregation perspective fully separates two challenges: the *epistemic* issue of identifying the beliefs (and confidence) that can or should be formed on the basis of a set of models; and the *pragmatic* question of their role in decision making. This is clear in Table 4: the resulting 'overall' decision rule depends not just on the aggregation parameters (first two columns), but also on the stakes, which under (16) regulate the degree of exhibited ambiguity aversion (Hill, 2013, 2019b). Indeed, this perspective can even bring to light hitherto unrecognised relationships between aggregation rules: for instance (Table 4, rows 2 & 4), SEU with geometric pooling and the average robust control rule turn out to be equivalent as concerns aggregation, differing only on the decision front. Moreover, separating out the aggregation part of the model misspecification challenge reconceptualises how to evaluate misspecification-sensitive decision models, and can suggest new directions.

[23] Hansen and Sargent's proof for multiplier preferences relies on the Lagrange multiplier theorem (Luenberger, 1969), and hence on the convexity of the constraint ($R(q\|m)$ in their case) as a function of q. The stated convexity conditions ensure the theorem applies to (16) and (17).

Model expertise Notably, confidence aggregation's ability to deal with expertise remains relevant when the 'experts' are models. In real applications, it is not uncommon 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 hurricanes; the latter in predicting ground-level temperature (Masson-Delmotte et al., 2021, Section 1.5.3). However, the distances typically underlying Bayesian Model Averaging and virtually all existing misspecification models (Table 4, rows 1–4) encode the assumption that all models are equally 'good'—they have comparable expertise—on all issues (Section 3.2). Like standard pooling rules (Section 3.1), existing model-misspecification decision rules thus cannot cope with intra-model cross-issue expertise-diversity.

As shown in Section 3, confidence aggregation can comfortably accommodate expertise differences within models. Using the same sort of expertise-sensitive confidence rankings (Definition 3) as in the derivation of our expertise-sensitive pooling rule (Definition 4) yields a expertise-sensitive generalisation of Bayesian Model Averaging (Table 4, row 5) at low stakes. At higher stakes levels, the same confidence ranking yields misspecification-sensitive decision rules that *can* accommodate differing degrees of expertise within models. Table 4 (rows 6–7) provides two examples, obtained by replacing the distance generating existing misspecification models by the expertise-sensitive distance introduced in Section 3.2. For the reasons set out in Section 3, they constitute arguably more pertinent misspecification-sensitive decision rules in situations where models' performance may vary across issues.

[24] As standard, the geometric pool involves a multiplicative constant, denoted χ.

[25] This generalised constraint rule takes $w(m) = 1$ for all $m \in \mathcal{M}$. Hansen and Sargent (2022) proposed the multiplier version of this model (i.e. (17) with $\otimes$ and ρ as specified in the Table and $w(m) = 1$ for all $m \in \mathcal{M}$); Cerreia-Vioglio et al. (2020, Theorem 1) axiomatise a version with general convex ρ, minimum $\otimes$ and convex, compact set of models $\mathcal{M}$. In such cases, the convexity assumptions needed to run Hansen and Sargent's (2008) argument hold (footnote 23), so these constraint preferences yield the same optima as the corresponding multiplier version, in the relevant decision problems.

[26] Although this representation concerns the case of finitely many models (Section 2) extension to a measure over (potentially infinitely many) models is straightforward. Hansen and Sargent (2007) used a multiplier version of this model (i.e. (17) with $\otimes$ and

ρ	$\otimes$	Stakes	'Overall' Decision rule
Reverse RE	Average	Low	$\mathbb{E}_{\sum_{\mathcal{M}} \frac{w(m)}{\sum_{m\in\mathcal{M}} w(m)} m} u(f)$ Bayesian Model Averaging & SEU (Steel, 2020)
RE	Average	Low	$\mathbb{E}_{\chi \prod_{\mathcal{M}} m^{\frac{w(m)}{\sum_{\mathcal{M}} w(m)}}} u(f)$ Geometric pooling & SEU (Dietrich, 2021)[24]
RE	Minimum	Medium / High	$\min_{\substack{q\in\Delta: \\ \min_{m\in\mathcal{M}} R(q\|m)\leqslant\eta}} \mathbb{E}_q u(f)$ 'Minimum' robust control (Hansen and Sargent, 2022; Cerreia-Vioglio et al., 2020)[25]
RE	Average	Medium / High	$\min_{\substack{q\in\Delta: \\ \sum_{\mathcal{M}} w(m)R(q\|m)\leqslant\eta}} \mathbb{E}_q u(f)$ 'Average' robust control (Hansen and Sargent, 2007; Lanzani, 2022)[26]
Exp-sens. RE	Average	Low	$\mathbb{E}_{\arg\min_{q\in\Delta} \sum_{\mathcal{M}} \sum_{j=1}^{l} w(m,l) R(q\|_{\mathcal{P}_j} \| m\|_{\mathcal{P}_j})} u(f)$ Expertise-sensitive pooling & SEU (Defn 4)
Exp-sens. RE	Minimum	Medium / High	$\min_{\substack{q\in\Delta: \\ \min_{\mathcal{M}} \sum_{j=1}^{l} w(m,l) R(q\|_{\mathcal{P}_j} \| m\|_{\mathcal{P}_j})\leqslant\eta}} \mathbb{E}_q u(f)$ Expertise-sensitive minimum robust control
Exp-sens. RE	Average	Medium / High	$\min_{\substack{q\in\Delta: \\ \sum_{\mathcal{M}} \sum_{j=1}^{l} w(m,l) R(q\|_{\mathcal{P}_j} \| m\|_{\mathcal{P}_j})\leqslant\eta}} \mathbb{E}_q u(f)$ Expertise-sensitive average robust control

Table 4: Confidence aggregation & confidence decision (Eq. (16)): Special cases.

Note: RE stands for 'Relative Entropy'; Exp-sens. RE is short for 'Expertise-sensitive Relative Entropy', i.e. ρ as in Definition 3 (and (5)), with the relative entropy classical distance. In the last three rows, $w : \mathcal{M} \times \{1, \ldots, l\} \to \Re_{\geqslant 0}$ is an assignment of weights to models and issues. Other notation is as in Sections 3, 4 and the text.

7 Discussion

This paper has largely focused on the contribution of confidence aggregation to recognised challenges for aggregating probability measures. Part of the related literature takes probabilities as primitive, rather than working

ρ as specified in the Table, for infinite $\mathcal{M}$), which was axiomatised by Lanzani (2022). In such cases, the convexity assumptions needed to run Hansen and Sargent's (2008) argument hold (footnote 23), so these constraint preferences yield the same optima as the corresponding multiplier version, in the relevant decision problems.

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: they fail to yield probability measures unless the weights are the same for all events, in which case one returns to standard linear pooling in the presence of a minimal Pareto-like condition (e.g. McConway, 1981; Genest and Zidek, 1986). This, and in particular the apparent impossibility in capturing within-person expertise diversity, has been argued to be a problem for linear pooling (e.g. French, 1985). The limit case in Example 3.4 (Section 3.3)—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 4, the expertise-sensitive pooling rule derived from confidence aggregation is always well-defined.

Confidence aggregation operates directly on confidence rankings, and hence, like pooling rules, does not require a preference setup to be applied. Just as pooling rules tacitly assume interpersonal comparison of probability judgements—one can say when two individuals are assigning the same probability—in direct application, confidence aggregation requires interpersonal comparison of confidence—one can tell when two individuals are talking about the same confidence level. Hill (2019a) discusses the problem of 'calibrating' confidence levels across individuals, providing and theoretically founding a scale for interpersonal confidence comparison which can be used for 'direct' applications of confidence aggregation. Interpersonal comparison can alternatively be provided by preferences, in the context of many decision models mentioned in Section 2.1 (see Section 4).

Another literature on belief aggregation works 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 the influential approach of Gilboa et al. (2004)—which characterises utilitarian aggregation of utility and linear pooling of probabilities—on these grounds. Several reactions in this literature work with preferences and consist in restricting the domain of the Pareto condition. Dietrich (2021) restricts it to cases where all agents have identical subjective probabilities, and adds a dynamical rationality condition of the sort discussed in Section 5. 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 approach to belief aggregation developed here, these make no attempt to retain the consensus-preservation intuition behind Pareto. By contrast, Bommier et al. (2021) present a condition preserving consensuses on prospects yielding identical distributions of outcomes for all individuals, and use it to provide a decision rule aggregating probabilistic beliefs. Under their procedure, the group 'belief' (distribution) used in the evaluation of a given prospect depends on the prospect in question, whereas ours produces a representation of group belief that is independent of the decision situation. Alon and Gayer (2016) and Stanca (2021) consider aggregation of SEU 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 (Section 4.1), confidence aggregation provides an aggregation rule for both sorts of non-expected utility 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, issue-wise Pareto. By contrast, the approach proposed here leverages the non-probabilistic structure of beliefs in aggregation, in concordance with the insight that confidence has a role in consensus formation (Introduction and Section 2.2). Nau (2002) proposes an aggregation rule for a confidence-based belief representation which is a special case of that used here (see 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. (3); see Nau (2002,

References

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

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

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

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

Bommier, A., Fabre, A., Goussebaïle, A., and Heyen, D. (2021). Disagreement aversion. *Available at SSRN 3964182*.

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

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. (2020). Making decisions under model misspecification. *arXiv preprint arXiv:2008.01071*.

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.

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. Publisher: Elsevier.

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

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

Figs 2 & 3).

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. (2007). Recursive robust estimation and control without commitment. *Journal of Economic Theory*, 136(1):1–27.

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. Publisher: Elsevier.

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.

Hill, B. (2022). Updating confidence in beliefs. *Journal of Economic Theory*, 199:105209.

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

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

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

Lanzani, G. (2022). Dynamic Concern for Misspecification.

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. Publisher: Cambridge University Press.

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

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

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. Publisher: Taylor & Francis.

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. Publisher: Elsevier.

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: Examples 3.2–3.4

In this Appendix, for completeness, we set out the details concerning the running Example in Sections 3.2–3.5, as well as some further comments.

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 $(p_L(L), p_R(R), p_B(B)) \in [0,1]^3$. Hence the inequalities defining the set $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$ (Eq. (9), Section 3.4) can be written in vector notation: $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$ is just the set of vectors $\mathbf{q} \in [0,1]^3$ satisfying the constraint $\mathbf{Aq} \leqslant \mathbf{r}$ 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 (6) 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:

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

where

$$\mathbf{p^L} = \begin{pmatrix} 0.9 \\ 0.1 \\ 0.09 \end{pmatrix}, \quad \mathbf{D^L} = \begin{pmatrix} 2w_L^L & 0 & 0 \\ 0 & 2w_R^L & 0 \\ 0 & 0 & 2w_B^L \end{pmatrix}$$

and similarly for Ray, with

$$\mathbf{p^R} = \begin{pmatrix} 0.1 \\ 0.9 \\ 0.09 \end{pmatrix}, \quad \mathbf{D^R} = \begin{pmatrix} 2w_L^R & 0 & 0 \\ 0 & 2w_R^R & 0 \\ 0 & 0 & 2w_B^R \end{pmatrix}$$

It follows that the minimisation problem (10) defining the centre of the confidence aggregation in Example 3.4 becomes the following simple quadratic optimisation problem under constraints:

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

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 solution is:

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

Otherwise, solving the minimisation problem yields:

$$\begin{aligned} p(L) =& \frac{\left(w_L^L p^L(L) + w_L^R p^R(L)\right) + \frac{w_B^L + w_B^R}{w_R^L + w_R^R} \left(w_L^L p^L(L) + w_L^R p^R(L) - w_R^L p^L(R) - w_R^R p^R(R)\right) + 1.09(w_B^L + w}{(w_L^L + w_L^R) + (w_B^L + w_B^R)(\frac{w_L^L + w_L^R}{w_R^L + w_R^R} + 1)} \\ p(R) =& \frac{\left(w_R^L p^L(R) + w_R^R p^R(R)\right) - \frac{w_B^L + w_B^R}{w_L^L + w_L^R} \left(w_L^L p^L(L) + w_L^R p^R(L) - w_R^L p^L(R) - w_R^R p^R(R)\right) + 1.09(w_B^L + w}{(w_R^L + w_R^R) + (w_B^L + w_B^R)(\frac{w_R^L + w_R^R}{w_L^L + w_L^R} + 1)} \\ p(B) =& p(L) + p(R) - 1 \end{aligned}$$

where, as specified, $p^L(L) = p^R(R) = 0.9$, $p^L(R) = p^R(L) = 0.1$ and $p^L(B) = p^R(B) = 0.09$. Taking limits as $\frac{w_B^L + w_B^R}{w_L^L + w_L^R} \to 0$ and $\frac{w_B^L + w_B^R}{w_R^L + w_R^R} \to 0$ and combining them with the expressions under slack constraints gives the expressions cited in Example 3.4.

Example 3.4 considers the case where $\frac{w_B^L+w_B^R}{w_L^L+w_L^R} \to 0$ and $\frac{w_B^L+w_B^R}{w_R^L+w_R^R} \to 0$, which, arguably, is closest to the example described in the Introduction. For completeness, note that, in the opposite case of $\frac{w_B^L+w_B^R}{w_L^L+w_L^R} \to \infty$ 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:

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

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 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.

B Characterising confidence aggregation: special cases

In this Appendix, we extend Theorem 1 to characterise, as special cases, confidence aggregation under the families of confidence-level aggregators mentioned in Section 2.2, as well as the following generalisation of the maximum aggregator.

Example B.1 (Generalised 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, is called a **generalised maximum aggregator**.

More specifically, we will provide results for the following stronger representation: $c^0 = \dot{F}_{\otimes}(c^1, \ldots, c^n)$, with $\dot{F}(c'^1, \ldots, c'^n) = \overline{\dot{\Phi}(c'^1, \ldots, c'^n)}$, where, for every $o \in O$ such that $\bigcup_{\mathbf{o}:\otimes\mathbf{o}=o} \bigcap_i c'^i(o_i) \neq \varnothing$

Axioms	Aggregator
Consensus Independence	Affine
Consensus Independence, Neutrality	Average
Consensus Join	Generalised Maximum
Consensus Join, Neutrality	Maximum
Consensus Meet, Neutrality	Minimum

Table 5: Characterisations of special case confidence-level aggregators, to be read in the context of Theorem B.1.

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

The only difference with respect to the representation involved in Theorem 1 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 (Section 2.2). It follows directly from the fact that confidence rankings are increasing in o that, if $c^0 = \dot{F}_{\otimes}(c^1, \ldots, c^n)$, then c^0 is a consensus-preserving confidence aggregation in the sense of Definition 2.

Recall that, under Assumptions 1 and 2, $\max\{D(\bullet), D(\bullet)\}$ is a monotonically increasing transformation of σ. 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 5.

Theorem B.1. *Suppose that O is infinite, let $\{\succ^i\}, \succ^0$ satisfy Assumptions 1 and 2 with $\max\{D(f), D(g)\} = \sigma(f, g)$ for all $f, g \in \mathcal{A}$. For each of the rows in Table 5: $\{\succ^i\}, \succ^0$ satisfy Corpus-wise Pareto, Consensus-based beliefs, Non-degeneracy and the axiom(s) in the first column of the table if and only if there exists a confidence-level aggregator $\otimes$ of the type specified in the second column such that $c^0 = \dot{F}_{\otimes}(c_1, \ldots, c_n)$, up to convex closure.*

We make no particular claim for any of the confidence-level aggregators in Table 5 on normative grounds; we present this result to illustrate the richness of the approach, and exemplify some simple aggregators.

Axiom (Consensus Independence). *For all tuples of stakes levels* $\mathbf{s_1}, \ldots, \mathbf{s_l}, \mathbf{t_1}, \ldots, \mathbf{t_m} \in \mathcal{S}^n$ *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* $\sum_{k=1}^{l} \alpha_k \mathbf{s}_l = \sum_{j=1}^{m} \beta_j \mathbf{t}_j$*, if, for some stakes levels* $s_1, \ldots, s_l, t_1, \ldots, t_m$*,* $\succ^0$ *does not respect the consensuses* $\succ_{\mathbf{s}_i}$ *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 any tuples of stakes levels* $\mathbf{s}, \mathbf{t}$ *exhibiting consensus, if* $\succ^0$ *respects the consensuses* $\succ_{\mathbf{s}}, \succ_{\mathbf{t}}$ *at* s*, then it respects the consensus* $\succ_{\mathbf{s} \vee \mathbf{t}}$ *at* s*.*

Axiom (Consensus Meet). *For any tuples of stakes levels* $\mathbf{s}, \mathbf{t}$ *exhibiting consensus, if* $\succ^0$ *respects the consensuses* $\succ_{\mathbf{s}}, \succ_{\mathbf{t}}$ *at* s*, then it respects the consensus* $\succ_{\mathbf{s} \wedge \mathbf{t}}$ *at* s*.*

Axiom (Neutrality). *For any stakes levels* s*, tuple of stakes levels* $\mathbf{s}$ *and permutation* π *such that* $\mathbf{s}, \pi(\mathbf{s})$ *exhibit consensus,* $\succ^0$ *respects the consensus* $\succ_{\mathbf{s}}$ *at* s *if and only if* $\succ^0$ *respects the consensus* $\succ_{\pi(\mathbf{s})}$ *at* s*.*

Where, for any $\mathbf{s}, \mathbf{t} \in \mathcal{S}^n$ and $\alpha \in [0,1]$, $(\alpha \mathbf{s} + (1-\alpha)\mathbf{t})_i = \alpha s_i + (1-\alpha) t_i$, $(\mathbf{s} \vee \mathbf{t})_i = \max\{s_i, t_i\}$ and $(\mathbf{s} \wedge \mathbf{t})_i = \min\{s_i, t_i\}$.

Figure 4: Axioms for special cases

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 and stakes level s with $\succ^0$ respecting the consensus $\succ_{\mathbf{s}}$ at s, we say that the consensus at s is *covered* when, for all acts f, g, if $f \not\succ_{\mathbf{s}} g$ then there exists a tuple $\mathbf{s}'$ exhibiting consensus with $\mathbf{s}' \not\geq \mathbf{s}$ such that $\succ^0$ respects the consensus $\succ_{\mathbf{s}'}$ at s and $f \not\succ_{\mathbf{s}'} g$. Otherwise, say that the consensus is *uncovered* at s. When the consensus $\succ_{\mathbf{s}}$ is covered, there is no f, g such that the absence of preference between them according to $\succ^0_s$ 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 $\succ^0$ 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 generalised 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, maximum and minimum confidence-level aggregators.

C Proofs

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

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

Proposition C.1. *c^0 is a consensus-preserving confidence aggregation of $(c^1, \dots, c^n)$ under $\otimes$ if and only if*

$$\iota^0(p) = \begin{cases} \otimes(\iota^1(p), \dots, \iota^n(p)) & \text{if } \forall i \ \iota^i(p) \in O \\ \varnothing & \text{otherwise} \end{cases}$$

where the implausibility function ι^i for c^i is as defined in Section 2.1.

Proof. By the definition, $p \in c^0(o)$ if and only if, for some $\mathbf{o}$ with $\otimes\mathbf{o} \leqslant o$, $p \in c^i(o_i)$ for all $i = 1, \ldots, n$. First consider p such that $\iota^i(p) \in O$ for all i. For such p, $p \in c^i(\iota^i(p))$ for all i and hence $p \in c^0(\otimes(\iota^1(p), \ldots, \iota^n(p)))$. Moreover, for any $\mathbf{o}$ with $\otimes\mathbf{o} < \otimes(\iota^1(p), \ldots, \iota^n(p))$, $o_i < \iota^i(p)$ for some i 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 $o' < \otimes(\iota^1(p), \ldots, \iota^n(p))$, $p \in c^0(o')$. The first clause of the required formula follows from the definition of ι. As concerns the other case, if there exists i with $\iota^i(p) = \varnothing$, then $p \notin c^i(o)$ for all $o \in O$, so, by the definition of confidence aggregation (notably Eq. (2)), $p \notin c^0(o)$ for all $o \in O$. So $\iota^0(p) = \varnothing$, as required by the second clause. □

We now prove Theorems 1 and B.1. Recall that, under Assumptions 1 and 2, $\{(c^i, D, u)\}, (c^0, D, u)$ denote the representations of the $\{\succ^i\}, \succ^0$. Moreover, as noted in Section 4.1 (footnote 21), $\max\{D(f), D(g)\} = \zeta \circ \sigma$ for some strictly increasing $\zeta : \mathbb{R} \to \mathbb{R}$. 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.

C.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, $X \neq \varnothing$. 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 C.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 C.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

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

Proof of Claim C.2. First consider $\succ_{\mathbf{s}}$ 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

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

Note that, since the c^i are closed and convex, so is their intersection. For every $f, g \in \mathcal{A}$, $f \succ_{\mathbf{s}} g$ if and only if $f \succ^i_{s_i} g$ for some i and $f \not\prec^i_{s_i} g$ for every i. By Assumption 1, this holds if and only if, 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 $\succ_{\mathbf{s}} = \succ_{\bigcap \mathbf{s}}$.

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 C.1. Note that if $G(X)$ is a finite set, then $\min G(X) \in G(X)$. The following proposition implies that this is the case when $G(X)$, and hence O, is infinite—and hence, given our assumptions, when the confidence rankings are upper semicontinuous.

Proposition C.2. *If the confidence rankings c^i are all upper semicontinuous, then, for any decreasing sequence $\mathbf{s}_j \in X$ with $\mathbf{s}^j \to \mathbf{s}$, $\mathbf{s} \in X$ and*

$G(\mathbf{s}) \leqslant \lim G(\mathbf{s}^j)$.

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, by the definition of G, $c^0(G(\mathbf{s})) \supseteq \bigcap_j \bigcap_{i=1}^n c^i(s_i^j)$, so $G(\mathbf{s}) \leqslant G(\mathbf{s}^j)$ for all j. Hence $G(\mathbf{s}) \leqslant \lim G(\mathbf{s}^j)$, as required. □

Claim C.3. For every $s \geqslant \min G(X)$, $\succ_s^0$ is represented by $\bigcup_{\mathbf{s}\in X: s \geqslant G(\mathbf{s})} \bigcap_i c^i(s_i)$ in the Bewley sense: i.e. for all $f, g \in \mathcal{A}$, $f \succ_s^0 g$ if and only if:

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

Proof. Fix a stakes level s with $s \geqslant \min G(X)$, and consider any $\mathbf{s}'$ with $G(\mathbf{s}') \leqslant s$. (By the previous observations guaranteeing the existence of a minimum, such $\mathbf{s}'$ exists.) By the definition of G, there exists $s'' \in X$ with $s'' \leqslant s$ and $\succ_{s''}^0 \subseteq \succ_{\mathbf{s}'}$. It follows from the nestedness properties of confidence rankings that $\succ_s^0 \subseteq \succ_{s''}^0 \subseteq \succ_{\mathbf{s}'}$. Since this holds for all $\mathbf{s}'$ with $G(\mathbf{s}') \leqslant s$, it follows that $\succ_s^0 \subseteq \bigcap_{\mathbf{s}\in X: s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$.

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 Corpus-wise Pareto, $f \succ_s^0 g$. So $\succ_s^0 \supseteq \bigcap_{\mathbf{s}\in X: s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$, and hence there is equality. It follows from Claim C.2 that (C.3) holds for all $s \geqslant \min G(X)$.

□

Since $c^0(s)$ represents $\succ_s^0$ 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 $\mathbf{s}, \mathbf{s}'$, 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

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

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 C.3 that (2) 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 $s < \min G(X)$, by the nestedness of confidence rankings, $\succ^0_s \subseteq \bigcup_{s'\in G(X)} \succ^0_{s'}$. However, by Consensus-based beliefs, 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 $\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i c^i(o_i) = \varnothing$, $c(o) = \bigcap_{s'\in G(X)} c^0(s') = c^0(\min G(X))$ (by the upper semicontinuity of confidence rankings), 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 (2), hence establishing the 'only if' direction.

The 'if' direction is a direct consequence of (2) and Claims C.1 and C.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 $G'(\mathbf{s}) = \otimes'(\mathbf{s})$. 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 (3) holds with equality in the place of the inequality.

C.1.2 Proof of Theorem B.1

Proof for Table 5, row 1 (affine aggregation). Let X be as defined in the proof of Theorem 1. Let

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

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 ith 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

$$L = \{(\mathbf{s}, s) \in Y \times \mathbb{R} : \mathbf{s} \in X_Y,\ \succ^0_s \nsubseteq \succ_{\mathbf{s}}\} = (X_Y \times \mathbb{R}) \backslash C_Y$$
$$U = \{(\mathbf{s}, s) \in C_Y : \exists s' \leqslant s, (\mathbf{s}, s') \in C_Y \backslash K_Y\}$$

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 C.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 C.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. By Proposition C.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 $X_Y \times \mathbb{R}$. 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 an 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 without loss of generality, ϕ, χ 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 ϕ. 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 C.6. For all $s \geqslant \inf G_{aff}(X)$, $\succ_s^0$ 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 $f, g \in \mathcal{A}$, $f \succ_s^0 g$ if and only if:

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

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 ϕ and the definition of NN, $\succ_s^0 \subseteq \succ_{\mathbf{s}}$. So $\succ_s^0 \subseteq \bigcap_{\mathbf{s}\in X:s=G_{aff}(\mathbf{s})} \succ_{\mathbf{s}}$.

We now establish the opposite containment. By Corpus-wise Pareto, $\succ_s^0 \supseteq \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C} \succ_{\mathbf{s}}$. 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 $\mathbf{s}'$, $\succ_s^0 \supseteq \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})\geqslant s} \succ_{\mathbf{s}} = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s} \succ_{\mathbf{s}}$, 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 $\succ_{\mathbf{s}'} \supseteq \bigcap_{\mathbf{s}'':(\mathbf{s}'',s)\in C,\ \mathbf{s}''\ngeqslant\mathbf{s}'} \succ_{\mathbf{s}''}$. So $\bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s} \succ_{\mathbf{s}} = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}} \cap \bigcap_{\mathbf{s}:(\mathbf{s},s)\in C,\ G_{aff}(\mathbf{s})>s,\mathbf{s}\ngeqslant\mathbf{s}'} \succ_{\mathbf{s}}$. 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 $\succ_s^0 \supseteq \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}}$.

So $\succ_s^0 = \bigcap_{\mathbf{s}:G_{aff}(\mathbf{s})=s} \succ_{\mathbf{s}}$; it follows from Claim C.2 that (C.4) holds for all $s \geqslant \inf G_{aff}(X)$. □

Since $c^0(s)$ represents $\succ_s^0$ 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

$$\otimes\mathbf{o} = \sum w_i o_i + \chi$$

Clearly, this is an affine confidence level aggregator. Moreover, by Claim C.6 and the fact that ζ 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, $s_k < w_i(\mathbf{s_k})_i + \chi$. 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 $\bigcap_i c^i((\mathbf{t_k})_i) \not\subseteq \bigcup_{\mathbf{s} \not\geq \mathbf{t_k},\ (\mathbf{s}, t_k) \in C} \bigcap_i c^i(s_i) \subseteq c^0(t_k)$. 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, $\sum_i w_i(\mathbf{t_k})_i + \chi \leq t_k$. So if $\succ_{\mathbf{t_k}}$ is an uncovered consensus at t_k for each k, it follows that $\sum_{k=1}^{m} \beta_k t_k \geq \sum_i w_i \sum_{k=1}^{m} \beta_k (\mathbf{t_k})_i + \chi$. 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 5, row 2 (averaging 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 $\geq$, 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 $s_j = s_k$, $G'_{aff}(\mathbf{s}) = \widehat{G_{aff}}(\mathbf{s}|_{\{1,\dots n\} \backslash \{k\}})$, 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 $\mathbf{s}$ with $s_j \neq s_k$ and $G'_{aff}(\mathbf{s}) = s$, since $(\mathbf{s}, s) \in C$, it follows that $(\hat{\mathbf{s}}, s) \in C$ by the specification of the case. So $\bigcap_i c^i(s_i) \subseteq \bigcap_i c^i(\hat{s}_i) \subseteq \bigcup_{\mathbf{s}' \in X : s = G'_{aff}(\mathbf{s}'),\ s_j = s_k} \bigcap_i c^i(s_i)$. Hence, for all $s \geqslant \inf G'_{aff}(X)$, $c^0(s) = \bigcup_{\mathbf{s} \in X : s = G'_{aff}(\mathbf{s})} \bigcap_{i=1}^n c^i(s_i)$, up to convex closure. So there is an affine aggregator representation with equal weights for j and k, as required.

Necessity of Neutrality is straightforward.

□

Proof for Table 5, row 3 (generalised maximum aggregator). Consider G as defined in the proof of Theorem 1. By Consensus Join, for any $\mathbf{s}, \mathbf{t}$, $G(\mathbf{s} \vee \mathbf{t}) \leqslant \max\{G(\mathbf{s}), G(\mathbf{t})\}$. However, by the monotonicity of G, since $\mathbf{s} \vee \mathbf{t} \geqslant \mathbf{s}, \mathbf{t}$, $G(\mathbf{s} \vee \mathbf{t}) \geqslant \max\{G(\mathbf{s}), G(\mathbf{t})\}$, so $G(\mathbf{s} \vee \mathbf{t}) = \max\{G(\mathbf{s}), G(\mathbf{t})\}$. For each $s \geqslant \min G(X)$, 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 $s_i > t^s_i$ 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 $s_i > t^s_i$. Hence $G(\mathbf{s}) < s$ if and only if there exists $s' < s$ with $s_i \leqslant t^{s'}_i$ 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 $\mathbf{s} \leqslant \mathbf{t}^{s'}$.

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

For $i = 1, \ldots, n$, define $\psi_i : O \to O$ by $\psi_i(o) = \zeta(\min\{s \in \mathcal{S} : \zeta(t^s_i) \geqslant o\})$. Since ζ is strictly increasing and, by the confidence representation, t^s_i is increasing in s for all i, ψ_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 ψ_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 5, row 4 & 5 (maximum & minimum aggregators). We present the proof for the maximum aggregator; the case of the minimum aggregator is similar. Consider $\mathbf{t^s}$, 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, $G((t_k^s)_j(t_j^s)_k\mathbf{t^s}) = G(\mathbf{t^s}) = s$; but since $t_j^s > t_k^s$, it follows by the properties of G established in the proof of part iii. that $G((t_k^s)_j(t_j^s)_k\mathbf{t^s}) > G(\mathbf{t^s}) = s$, which is a contradiction. So $t_j^s = t_k^s$ for all j, k and s. Hence, for ψ_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.

□

C.2 Proofs of results in Section 3

Proof of Proposition 1. By (3), 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 3, $\iota^i(p) = w^i \sum_{\omega\in\Omega}(p(\omega) - p_i(\omega))^2$, 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 Δ.

For the second row of the Table, $\iota^i(p) = w^i R(p\|p_i)$, so the centre of c is $p = \arg\min \sum_{i=1}^{n} w^i R(p\|p_i)$. Yet:

$$\begin{aligned}
\sum_{i=1}^{n} w^i R(p\|p_i) &= -\sum_{i=1}^{n} w^i \sum_{\omega\in\Omega} p(\omega)(\log \frac{p_i(\omega)}{p(\omega)}) \\
&= -\sum_{\omega\in\Omega} p(\omega) \log \left(\prod_{i=1}^{n} \frac{p_i(\omega)^{w^i}}{p(\omega)^{w^i}} \right) \\
&= -\left(\sum_{i=1}^{n} w^i\right) \sum_{\omega\in\Omega} p(\omega) \log \left(\frac{\prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}}{p(\omega)} \right) \\
&= \left(\sum_{i=1}^{n} w^i\right) \left[-\sum_{\omega\in\Omega} p(\omega) \log \left(\frac{\prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}}{p(\omega)} . \frac{1}{\sum_{\omega\in\Omega} \prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}} \right) \right. \\
&\quad \left. + \log \left(\frac{1}{\sum_{\omega\in\Omega} \prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}} \right) \right] \\
&= \left(\sum_{i=1}^{n} w^i\right) \left[-\sum_{\omega\in\Omega} p(\omega) \log \left(\frac{GM(p_i)(\omega)}{p(\omega)} \right) + \log \left(\frac{1}{\sum_{\omega\in\Omega} \prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}} \right) \right] \\
&= \left(\sum_{i=1}^{n} w^i\right) \left[R(p\|GM(p_i)) + \log \left(\frac{1}{\sum_{\omega\in\Omega} \prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}} \right) \right]
\end{aligned}$$

where $GM(p_i)(\omega) = \frac{\prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}}{\sum_{\omega\in\Omega} \prod_{i=1}^{n} p_i(\omega)^{\frac{w^i}{\sum_{i=1}^{n} w^i}}}$. This expression is clearly minimised at $p = GM(p_i) \in \Delta$, so the centre of c is $GM(p_i)$, as required.

For the third row, $\iota^i(p) = w^i R(p_i\|p)$, so the centre of c is $p = \arg\min \sum_{i=1}^{n} w^i R(p_i\|p)$. Yet:

$$\begin{aligned}
\sum_{i=1}^{n} w^i R(p_i \| p) &= -\sum_{i=1}^{n} w^i \sum_{\omega\in\Omega} p_i(\omega)(\log \frac{p(\omega)}{p_i(\omega)}) \\
&= \sum_{i=1}^{n} w^i \sum_{\omega\in\Omega} p_i(\omega)\log p_i(\omega) - \sum_{\omega\in\Omega} \log p(\omega) \sum_{i=1}^{n} w^i p_i(\omega) \\
&= \sum_{i=1}^{n} w^i \sum_{\omega} p_i(\omega)\log p_i(\omega) - \left(\sum_{i=1}^{n} w^i\right)\left(\sum_{\omega\in\Omega} AM(p_i)(\omega)\log AM(p_i)(\omega)\right) \\
&+ \left(\sum_{i=1}^{n} w^i\right)\left(\sum_{\omega\in\Omega}(\log AM(p_i)(\omega) - \log p(\omega))AM(p_i)(\omega)\right) \\
&= \sum_{i=1}^{n} w^i \sum_{\omega} p_i(\omega)\log p_i(\omega) - \left(\sum_{i=1}^{n} w^i\right)\left(\sum_{\omega} AM(p_i)(\omega)\log AM(p_i)(\omega)\right) \\
&+ \left(\sum_{i=1}^{n} w^i\right) R(AM(p_i)\|p)
\end{aligned}$$

where $AM(p_i) = \sum_{i=1}^{n} \frac{w^i}{\sum_{i=1}^{n} w^i} p_i$. This expression is clearly minimised at $p = AM(p_i) \in \Delta$, so the centre of c is $AM(p_i)$, as required.

□

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

Proposition C.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^{RE} be the w^L reverse relative entropy confidence ranking generated by p^L. 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 $c^{RE}(o) \subseteq \mathcal{L}_\epsilon$ if and only if $c^{RE}(o) \subseteq \mathcal{R}_\epsilon$.*

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

Both the distance functions involved (Euclidean distance, relative entropy) are functions 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, in the Euclidean case,

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

In the relative entropy case,

$$
\begin{aligned}
&\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p(\omega_{LR}), p(\omega_L), p(\omega_N)) \\
=& - p^L(\omega_{LR}) \log\left(\frac{p(\omega_{LR})}{p^L(\omega_{LR})}\right) - p^L(\omega_L) \log\left(\frac{p(\omega_L)}{p^L(\omega_L)}\right) - p^L(\omega_N) \log\left(\frac{p(\omega_N)}{p^L(\omega_N)}\right) \\
&- (1 - p^L(\omega_{LR}) - p^L(\omega_L) - p^L(\omega_N)) \log\left(\frac{(1 - p(\omega_{LR}) - p(\omega_L) - p(\omega_N))}{(1 - p^L(\omega_{LR}) - p^L(\omega_L) - p^L(\omega_N))}\right)
\end{aligned}
$$

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. □

Proof of Proposition 2. Take $o = 2\left(w_L^L \epsilon^2 + w_B^L (\max\{\epsilon - 0.81, 0\})^2\right)$. q, defined by $q(\omega_{LR}) = p^L(\omega_{LR}) - \max\{\epsilon - 0.81, 0\} = 0.09 - \max\{\epsilon - 0.81, 0\}$, $q(\omega_R) = p^L(\omega_R) + \max\{\epsilon - 0.81, 0\} = 0.01 + \max\{\epsilon - 0.81, 0\}$, $q(\omega_L) = p^L(\omega_L) - \epsilon = 0.81 - \min\{\epsilon, 0.81\}$ and $q(\omega_N) = p^L(\omega_N) + \min\{\epsilon, 0.81\} = 0.09 + \min\{\epsilon, 0.81\}$ is a probability measure over Ω. Moreover, $\sum_{j=\{L,R,B\}} 2w_j^L (q(j) - p^L(j))^2 = 2\left(w_L^L \epsilon^2 + w_B^L (\max\{\epsilon - 0.81, 0\})^2\right)$, so $q \in c^L(o)$. Since, for any q' with $q'(L) < 0.9 - \epsilon$, $\sum_{j=\{L,R,B\}} 2w_j^L (q'(j) -$

$p^L(j))^2 > 2\left(w_L^L\epsilon^2 + w_B^L\left(\max\{\epsilon - 0.81, 0\}\right)^2\right)$, such $q' \notin c^L(o)$, so $c^L(o) \subseteq \mathcal{L}_\epsilon$. For any $\delta \in [0, 0.9]$, consider q_δ 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. $\sum_{j=\{L,R,B\}} 2w_j^L(q_\delta(j) - p^L(j))^2 = 2\left(w_R^L\delta^2 + w_B^L\left(\max\{0, \delta - 0.09\}\right)^2\right)$. Noting that $w_R^L\epsilon^2 + w_B^L\left(\max\{0, \epsilon - 0.09\}\right)^2 < w_L^L\epsilon^2 + w_B^L\left(\max\{\epsilon - 0.81, 0\}\right)^2$ if and only if $w_B^L\frac{1}{\epsilon^2}\left(\left(\max\{0, \epsilon - 0.09\}\right)^2 - \left(\max\{\epsilon - 0.81, 0\}\right)^2\right) < w_L^L - w_R^L$, 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 $\sum_{j=\{L,R,B\}} 2w_j^L(q_\delta(j) - p^L(j))^2 \leqslant 2\left(w_L^L\epsilon^2 + w_B^L\left(\max\{\epsilon - 0.81, 0\}\right)^2\right) = o$, so $c^L(o) \nsubseteq \mathcal{R}_\epsilon$, as required.

□

C.3 Proofs of results in Section 5

Proof of Theorem 2. Fix E and ρ_E, and define $c^\rho : O \to 2^\Delta \backslash \varnothing$ by $c^\rho(o) = \{p \in \Delta : p(E) \geqslant \rho_E(o)\}$. Clearly, for any confidence ranking c, $c|\rho_E = \overline{\Phi}$ for $\Phi(o) = (c(o) \cap c^\rho(o))_E$, whenever $c(o) \cap c^\rho(o) \neq \varnothing$ (and it is undefined otherwise).

By Definition 6 and the definition of $F_\otimes$, for every $o \in O$ such that $\left(\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i c^i(o_i)\right) \cap c^\rho(o) \neq \varnothing$

$$\begin{aligned}
&F_\otimes(c_1, \ldots, c_n)|\rho_E(o)\\
&= \left(\left(\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i c^i(o_i)\right) \cap c^\rho(o)\right)_E\\
&= \left(\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i \left(c^i(o_i) \cap c^\rho(o)\right)\right)_E\\
&= \left(\bigcup_{\mathbf{o}:\otimes\mathbf{o}\leqslant o}\bigcap_i \left(c^i(o_i) \cap c^\rho(o)\right)_E\right)\\
&= F_\otimes(c_1|\rho_E, \ldots, c_1|\rho_E)(o)
\end{aligned}$$

as required. □