Confidence, consensus and aggregation*

Brian Hill
GREGHEC, HEC Paris & CNRS†

May 26, 2023

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 characterised by a variant of the Pareto condition that enjoins respecting consensuses borne of compromise. Confidence aggregation recoups standard probability aggregation rules, such as linear pooling, as special cases, whilst avoiding the spurious unanimity issues that have plagued such rules. Moreover, it generates a new family of probability aggregation rules that can faithfully accommodate within-person expertise diversity, hence resolving a longstanding challenge. Confidence aggregation also outperforms linear aggregation: the group beliefs it provides are closer to the truth, in expectation. Finally, confidence aggregation is dynamically rational: it commutes with update.

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

JEL codes: D70, D81.

* Acknowledgements and thanks to be added.

† 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 the economic, statistics, psychology and risk analysis literatures is linear pooling, which takes a weighted average of probabilistic beliefs. 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 come under increasing criticism.

One central problem with linear pooling as a belief aggregation mechanism 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. Whilst 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,

	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

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.

the agreement on this probability is *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 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, 2017; 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 some competence on the real estate sector but more expertise on the labour market (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, 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, including information, other beliefs, or reasons (Mongin and Pivato, 2020; Dietrich, 2021; Bommier et al., 2021). One could add evidence to the list: presumably Laura's and Ray's similar judgements on the 'Both' proposition are based on different evidence, supporting the low probabilities they assign to Real-Estate and Labour respectively. This suggests that an agent's declared probability concerning an event does not exhaust her relevant judgements pertaining to that event. As such, it echoes a position defended in the literature on rational belief representation, decision and learning: a probability judgement does not fully capture all relevant aspects of a rational agent's state of belief concerning an event. For instance, Hill (2013, 2019b) has argued that a rational belief state also comprises of the agent's *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 some opinions to retain the possibility of others. They may 'put aside' some beliefs to focus on 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 in any such consensus be preserved in the group's beliefs, or at least that they not be ruled out. In other words, the probability judgements that should be preserved are those *that belong to corpus-level consensuses*—in the sense that they fit into coherent sets of consensus positions encompassing all the relevant issues.

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 linear pooling. We first show that popular probabilistic opinion pooling rules can be reproduced as special cases of confidence aggregation, corresponding to particular assumptions on individuals' confidence in their beliefs. This sheds light on the comparison with existing approaches: whereas classic pooling rules are essentially based on assumptions about what individuals are willing to compromise to arrive at group beliefs, our approach uses precisely the compromises provided by the individuals themselves, as encoded in the confidence they have in their beliefs.

This analysis also sets the stage for the integration of expertise diversity across issues within individuals. 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 than the latter. 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.

In an application of our approach, 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.

Our third contribution is to develop a preliminary comparative appraisal of the *performance* of confidence aggregation, as measured by how correct the resultant group beliefs are in expectation. Confidence is a topic of intense research in Cognitive Psychology, where a significant literature has shown that, in many domains, a person's confidence in her judgement correlates positively with the judgement's correctness (Koriat, 2012a; Fleming and Lau, 2014; Rahnev et al., 2020). Incorporating such findings into a toy model of the cognitive processes underlying confidence formation, we show that confidence aggregation leads, in expectation, to group beliefs that are more correct than those provided by linear pooling, according to standard measures such as the Brier score. The essential insight is that confidence assessments convey useful information about the chances that a probability judgement is correct and, while confidence aggregation mobilises this information, classic pooling rules ignore it, tacitly imposing their own 'automatic' concept of confidence.

In a final contribution, we briefly consider the issue of dynamic rationality, which is typically evoked to justify the geometric pooling rule (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.

The paper is structured in a modular fashion: apart from Section 2, which sets out the framework and the aggregation rule, the sections can be read in any order or skipped at will. Section 3 shows how confidence aggregation overcomes the challenges to linear pooling, and contains a serious of examples that illustrate the approach. In particular, Section 3.1 establishes that standard pooling rules can be recovered as a special case of confidence aggregation, whereas Sections 3.2 and 3.3 illustrate how it can faithfully accommodate within-person expertise diversity. Section 3.3 also develops 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. Section 5 compares the performance of confidence aggregation to that of linear pooling, and Section 6 considers the topic of dynamic rationality. Section 7 discusses remaining related literature and outstanding issues, including consequences of our approach for aggregation of and robust decision with models. Proofs and supplementary material are contained in the Appendices.

2 Confidence aggregation

2.1 Setup: Beliefs and Confidence

Let Ω be a non-empty set of states. For the purposes of exposition, Ω can be taken to be finite, though extension to the infinite case is straightforward. Subsets of Ω are called events and Δ is the set of probability measures over Ω. $O \subseteq \mathbb{R}$ 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 an interval in $\mathbb{R}$, with the associated topology.[1] In the infinite case, all functions mentioned will be assumed to be continuous. 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 $i = 1, \ldots, n$.

We adopt the representation of confidence in beliefs proposed by Hill (2013); see Hill

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

(2019b) for a detailed defence of this approach as particularly appropriate for normative applications. The belief state of an agent—incorporating confidence—is represented by a *confidence ranking*: a function $c : O \to 2^{\Delta} \backslash \varnothing$ that is increasing in the containment order on sets.[2] For each confidence level o, $c(o)$ is the (non-empty) set of priors representing the beliefs held with confidence at least o. If O is infinite, we assume that c has the following upper semicontinuity property: for any decreasing sequence $o_i \in O$ with $o_i \to o$, $c(o) = \bigcap_i c(o_i)$. Hill (2013, 2016) axiomatise preferences involving confidence rankings satisfying a stronger continuity property. For any $o \in O$ and 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.

A confidence ranking c is *centred* if, for some $o \in O$, $c(o)$ is a singleton. By the monotonicity property of confidence rankings, if c is centred, there is a unique $p \in \Delta$ such that $c(o) = \{p\}$ for some o; we call this the *centre* of the confidence ranking. As discussed in Hill (2013), centred confidence rankings represent Bayesians with confidence: agents who assign a precise probability to every event (namely, that given by p), though may have more confidence in some judgements than others. A confidence ranking c is *convex* if, for every $o \in O$, $c(o)$ is a convex set; it is *closed* if, for every $o \in O$, $c(o)$ is a 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.

The previous definition of confidence rankings generates two alternative equivalent representations. Firstly, note that each probability judgement—judgements 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 measures where the judgement holds. Noting this, the function $conf : 2^{\Delta} \to O \cap \{\varnothing\}$, defined by:

$$conf(\mathcal{P}) = \begin{cases} \varnothing & \bigcap_{o \in O} c(o) \nsubseteq \mathcal{P} \\ \max\{o : c(o) \subseteq \mathcal{P}\} & o/w \end{cases} \tag{1}$$

picks out, for any probability judgement $\mathcal{P}$, the agent's confidence in $\mathcal{P}$—the largest confidence level at which $\mathcal{P}$ is held if it is held, and nothing otherwise. The confidence representation also generates a unique *implausibility function* $\iota : \Delta \to O \cap \varnothing$ defined by $\iota(p) = \min\{o \in O : p \in c(o)\}$ whenever the set is non-empty, and $\iota(p) = \varnothing$ otherwise. This yields the 'implausibility' of each probability measure, in terms of the smallest confidence level such that the probability measure doesn't contradict the judgements held with that much confidence.[3]

[2] I.e. for all $o \geqslant o'$, $c(o) \supseteq c(o')$.

[3] Note that c can be defined from ι: $c(o) = \{p \in \Delta : \iota(p) \leqslant o\}$.

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 one 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.

2.2 Consensus and confidence aggregation

To introduce our 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' respective 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 by $\mathbf{o}$.[4]

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 higher: 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 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 unwilling to compromise.

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.

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

In our preference-based characterisation, 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 (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**.

Example 2.4 (Maximum aggregator)**.** The special case of the generalised maximum aggregator with the same transformation for all individuals is the **maximum aggregator**, defined by $\otimes\mathbf{o} = \psi(\max\{o_i\})$, where $\psi : O \to O$ is as above.

We can now introduce our confidence aggregation rule. Since, in the general setup used in this paper, agents' 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 $F_\otimes$, taking profiles of confidence rankings into confidence rankings, 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-preservation aggregation for a fixed confidence ranking c^0 and profile $(c^1, \ldots, c^n)$.

Definition 1. 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 $\otimes$. More specifically, it 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 a judgement, neither judgement will be retained in the group beliefs with confidence o.

A noteworthy consequence of this aggregation procedure is that group and individual confidence in a judgement co-vary. More precisely, because of the monotonicity of $\otimes$, the group confidence in a judgement is higher when the individual beliefs feeding into the relevant consensuses are held at higher confidence levels. This appears to be a reasonable property for a procedure for aggregating beliefs and confidence.

The use of consensuses corresponding to confidence levels less than and equal to o ensures that c^0 is a well-defined confidence ranking, without requiring any assumptions on $\otimes$. As discussed in Appendix A, 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. Note finally that, by the definitions above, the previous notion can be formulated in terms of implausibility functions (Proposition 3, Appendix B.1): c^0 is a consensus-preserving aggregation of $(c^1, \ldots, c^n)$ under $\otimes$ if and only if, for all $p \in \Delta$

$$\iota^0(p) = \otimes(\iota^1(p), \ldots, \iota^n(p)) \tag{3}$$

3 Confidence, probability aggregation and expertise diversity

In this section, we consider the challenges to the linear pooling rule for probability aggregation, showing that our confidence aggregation proposal is not only capable of surmounting them, but does so very naturally. We first show that standard probability aggregation rules can be recovered as special cases of confidence aggregation. Moreover, they correspond to particular assumptions about individuals' confidence in their probability judgements. We then consider applications of confidence aggregation that do not make such assumptions, showing that, once one frees oneself from them, the rule naturally resolves both the spurious unanimity and the within-person cross-issue expertise diversity challenges. This discussion also contains several examples of confidence aggregation, which provide illustrations of the approach.

3.1 Recovering probability aggregation from confidence aggregation

Probability aggregation takes as input a profile of probability measures $\mathbf{p} = (p_1, \ldots, p_n)$. To connect pooling rules operating on such profiles with confidence aggregation, recall that each centred confidence ranking is naturally associated with a unique probability measure, namely its centre (Section 2.1). Conversely, given a probability measure, there are several ways of generating a centred confidence ranking with the specified probability measure as its centre. The definition below provides some examples. Here and throughout this section, we take $O = [0,1]$.

Definition 2. Let $p \in \Delta$ be a probability measure, $\omega' \in \Omega$ and $\Omega' = \Omega \backslash \{\omega'\}$. Then, for every $w \in (0,1)$:

1. If the state space Ω is finite, the w Euclidean confidence ranking generated by p is defined by $c(o) = \{q \in \Delta : w \sum_{\omega \in \Omega'} (q(\omega) - p(\omega))^2 \leqslant o\}$.

2. The w relative entropy confidence ranking generated by p is defined by $c(o) = \{q \in \Delta : wR(q\|p) \leqslant o\}$, where $R(q\|p) = -\sum q(\omega)(\log \frac{p(\omega)}{q(\omega)})$ is the Kullback-Leibler divergence, or relative entropy.

3. The w reverse relative entropy confidence ranking generated by p is defined by $c(o) = \{q \in \Delta : wR(p\|q) \leqslant o\}$, where $R(p\|q) = -\sum p(\omega)(\log \frac{q(\omega)}{p(\omega)})$ is the relative entropy.[5]

These examples hardly exhaust the ways of generating centred confidence rankings from probability measures. They do illustrate an important way of doing so: by taking the sets in the confidence ranking to be those which are closer to the specified probability measure under some distance on Δ (be it a metric, as in the first case, or a divergence as in the other cases).

These provide 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 one of the generation procedures just illustrated. 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, so the procedure yields a unique probability measure. This schema is summarised in Figure 1.

The following result compares this method of aggregating probability measures to standard pooling rules.

Theorem 1. *Let $\mathbf{p} = (p_1, \ldots, p_n) \in \Delta^n$ be a profile of probability measures. Then, for every n-tuple of weights $(w^1, \ldots, w^n)$, with $w^i \geqslant 0$ for all i:*

[5] Recall that the relative entropy is not symmetric.

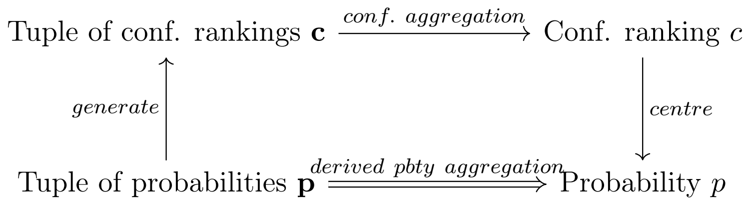

Figure 1: Using confidence aggregation to generate probability aggregation rules

1. *if c is the consensus-preserving confidence aggregation under an average confidence-level aggregator of w^i Euclidean confidence rankings generated by p_i (for each i), then its centre p is the linear pool of p_i, with weights $\frac{w^i}{\sum_{i=1}^n w^i}$: $p = \sum \frac{w^i}{\sum_{i=1}^n w^i} p_i$;*

2. *if c is the consensus-preserving confidence aggregation under an average confidence-level aggregator of w^i relative entropy confidence rankings generated by p_i (for each i), then its centre p is the geometric pool of p_i, with weights $\frac{w^i}{\sum_{i=1}^n w^i}$: $p(\omega) \propto \prod p_i^{\frac{w^i}{\sum_{i=1}^n w^i}}(\omega)$ for all $\omega \in \Omega$;*

3. *if c is the consensus-preserving confidence aggregation under an average confidence-level aggregator of w^i reverse relative entropy confidence rankings generated by p_i (for each i), then its centre p is the linear pool of p_i, with weights $\frac{w^i}{\sum_{i=1}^n w^i}$: $p = \sum_{i=1}^n \frac{w^i}{\sum_{i=1}^n w^i} p_i$;*

Hence the two most prominent pooling rules in the probability aggregation literature 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 this result is the use of specific confidence rankings for the individuals in the group. As is clear from the comparison of the clauses in the Theorem, 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 w^i Euclidean or w^i reverse relative entropy.[6]

This observation brings a new perspective on the evaluation of these pooling rules. Assessing the normative credentials of linear pooling, for instance, is equivalent to adopting the confidence framework and appraising the 'rationality' of the w^i Euclidean or w^i

[6] Given that, as noted above, a distance and a probability measure generate a confidence ranking, Theorem 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 agents' belief states—and, as shall be clear below, evaluate them as such.

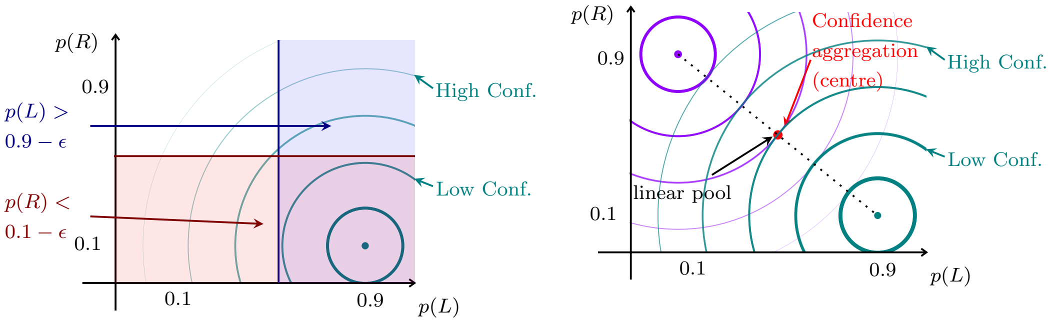

(a) Illustration of Proposition 5.
Note: The blue area represents the probability judgement that $p(L)$ is within ϵ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement that $p(R)$ is within ϵ of $p^L(R) = 0.1$. The confidence in these judgements (corresponding to the largest circular set contained in each area; Section 2.1) is the same.

(b) Illustration of Theorem 1.
Note: The red point is the centre of the result of confidence aggregation applied to the two confidence rankings (Theorem 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 Theorem 1.
Note: Each graph shows the space of pairs of probability values $(p(L), p(R))$ for the Labour and Real Estate events (L and R; Example 3.1). The areas (sets of probability values) enclosed by the green circles represent the w^L Euclidean confidence ranking generated by Laura's probability p^L (Definition 2): they are the projection of the confidence ranking into this space. Larger, lighter circles correspond to higher confidence levels. The purple circles represent the w^R Euclidean confidence ranking generated by p^R (Ray's probabilities), with $w^R = w^L$.

reverse relative entropy methods of generating confidence rankings. More importantly, it suggests a strategy for developing probability aggregation rules that overcomes the challenges cited in the Introduction. If the weaknesses of linear pooling are in fact connected to how rankings are generated from single probability measures, then applying confidence aggregation with different generation methods may produce pooling rules that avoid these problems. We now show that this strategy can be used to overcome both the spurious unanimity and within-person expertise diversity challenges.

3.2 Representing expertise using confidence rankings

One specificity of the confidence rankings involved in Theorem 1, related to the fact that they are based on distances in probability space, 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 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 point on the example from the Introduction.

Example 3.1. To formalise the example from the Introduction, consider a four-state state space $\Omega = \{\omega_{LR}, \omega_L, \omega_R, \omega_N\}$ where ω_{LR} (respectively, ω_L, ω_R, ω_N) is the state where 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 $L = \{\omega_{LR}, \omega_L\}$ and the corresponding event concerning the real-estate sector is $R = \{\omega_{LR}, \omega_R\}$. Consider Laura, whose probability judgements 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 that the probability of L is greater than $0.9-\epsilon$, and the judgement that the probability of R is less than $0.1+\epsilon$. Note that these judgements involve moving the same 'distance' from her best-guess probability judgement for L (0.9) and R (0.1) respectively. Which of the judgements is she more confident in?

Proposition 5 (Appendix B.2) shows that, under the two confidence ranking generating procedures yielding linear pooling—w^L Euclidean and w^L reverse relative entropy confidence rankings—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 who have the same confidence in the probability judgements encoded in their probability measure p^i, no matter the issues that the judgements concern. As such, these rankings fail to properly capture an individual who has different confidence in judgements pertaining to different issues. The previous example is arguably such a case. As discussed in the Introduction, Laura has more expertise on one issue (labour) than another (realestate). But an expertise difference typically translates into a difference in confidence: *ceteris paribus* she will have more confidence in her judgements concerning the issue of her expertise than in those that do not. The confidence rankings discussed above, based on (classic) distances on the probability space, assume that there is no within-person cross-issue difference in expertise.

This observation, combined with Theorem 1, 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 expertise differences. It thus suggests that confidence aggregation applied to confidence rankings that *do* correctly capture expertise differences could incorporate more faithfully such differences into the group beliefs. We now confirm this suggestion, and show how it can produce new expertise-sensitive probabilistic belief aggregation rules.

For presentation, we focus on issues that can be related to events in Ω; see Example 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.

Recall that the set of probability distributions on Ω is $\Delta = \Delta(\Omega)$; for any partition $\mathcal{P}_j$ of Ω, $\Delta(\mathcal{P}_j)$ is the set of probability distributions on $\mathcal{P}_j$. For any $p \in \Delta$ and partition $\mathcal{P}_j$, let $p|_{\mathcal{P}_j} \in \Delta(\mathcal{P}_j)$ be the projection of p into $\Delta(\mathcal{P}_j)$. For a set 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. As shall be illustrated shortly, this set is typically defined by a set of inequalities. 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. Consider the example from the Introduction, with the state space and events defined in Example 3.1. Each of the three issues mentioned in the example corresponds to a two-element partitions: $\mathcal{P}_L = \{L, L^c\}$ (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).

Each probability measure p over $\{L, L^c\}$ is determined by $p(L)$, and similarly for the other partitions. So each tuple $(p_L, p_R, p_B) \in \Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ is fully characterised by the vector $(p_L(L), p_R(R), p_B(B)) \in [0,1]^3$. The set $P_{\mathcal{P_L},\mathcal{P_R},\mathcal{P_B}}$ is defined by the following linear inequalities imposed by the fact that $B = L \cap R$: for any $p \in \Delta$

$$\begin{aligned} p(L) &\geqslant p(B) \\ p(R) &\geqslant p(B) \\ 1 &\geqslant p(L) + p(R) - p(B) \end{aligned} \tag{4}$$

Using the vector notation, justified as observed above, $P_{\mathcal{P_L},\mathcal{P_R},\mathcal{P_B}}$ is just the set of vectors

$\mathbf{q} \in [0,1]^3$ satisfying the constraint $\mathbf{Aq} \leqslant \mathbf{r}$ where

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

Clearly each $\mathbf{q}$ satisfying this constraint determines a unique probability measure on $\Omega = \{\omega_{LR}, \omega_L, \omega_R, \omega_N\}$, so the sequence of partitions $\mathcal{P}_L, \mathcal{P}_R, \mathcal{P}_B$ is rich.

A (statistical) *distance* d is the specification, for each partition $\mathcal{P}$ (including Ω itself), of a function $d : \Delta(\mathcal{P})^2 \to [0, \infty]$ such that: $d(q,p) = 0$ if and only if $p = q$; and $d(\bullet, q)$ is a lower semicontinuous function, for all $q \in \Delta(\mathcal{P})$.[7] Clearly, many metrics (like the Euclidean metric with Ω finite) and divergences (such as relative entropy) are distances. A distance d is *convex* if, for every $\mathcal{P}$ and $p \in \Delta(\mathcal{P})$, the function $d(\bullet, p)$ is strictly convex.[8] It is straightforward to check that the Euclidean distance, as well as the two distances defined from the relative entropy—$d_1(q,p) = R(q\|p)$ and $d_2(q,p) = R(p\|q)$—are convex.

We now define a family of centred confidence rankings that can capture cross-issue expertise diversity.

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

$$c(o) = \left\{ q \in \Delta : \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \leqslant o \right\} \tag{5}$$

For such confidence rankings, at each confidence level, the 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. Note that the confidence rankings in Definition 2 are special cases of $\mathbf{w}$ d-confidence rankings involving a single partition $\mathcal{P} = \Omega$ and particular choices of distance d.

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 our example.

Example 3.3. Consider p^L as defined in Example 3.1, and suppose that Laura's confidence ranking is generated by it with Euclidean distance and vector of weights $\mathbf{w^L} =$

[7] Throughout, we take the weak* topology on $\Delta(\mathcal{P})$ (Aliprantis and Border, 2007). Note that one can imagine conditions relating the d functions across partitions; no such conditions are required for the developments below.

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

(w_L^L, w_R^L, w_B^L). I.e. Laura has the confidence ranking:

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

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*, the agent is represented as having more confidence in judgements concerning issues with higher weights. This is borne out by the following proposition.

Proposition 1. *Suppose that $w_L^L > w_R^L$ and $0.8w_B^L < w_L^L - w_R^L$. Then, for every $\epsilon \in [0, 0.9]$, and $\mathcal{L} = \{p \in \Delta : p(L) \geqslant 0.9 - \epsilon\}$, $\mathcal{R} = \{p \in \Delta : p(R) \leqslant 0.1 + \epsilon\}$, there exists $o \in O$ with $c^L(o) \subseteq \mathcal{L}$ but $c^L(o) \nsubseteq \mathcal{R}$.*

Whenever w_B^L is not too large, if $w_L^L > w_R^L$, then any judgement 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. Figure 3a illustrates the intuition: when $w_L^L > w_R^L$, the sets in the confidence ranking have an 'elliptical' shape which is thinner along the L dimension, hence translating higher confidence in judgements on this issue. So $w_L^L > w_R^L$ reflects higher confidence *ceteris paribus* in judgements about the labour market as compared to the real estate sector, and would be a natural assumption for Laura's confidence ranking, given her expertise. The clause concerning w_B^L is related to the constraints that a given value of $p(B)$ places on the possible values of $p(L)$ and $p(R)$, as will be discussed 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 (where where $p^R(\omega_{LR}) = 0.09,\ p^R(\omega_L) = 0.01,\ p^R(\omega_R) = 0.81, p^R(\omega_N) = 0.09$), with weights $\mathbf{w^R} = (w_L^R, w_R^R, w_B^R)$ where $w_R^R > w_L^R$, translating his relative expertise in real estate.

Remark 1. For future reference, it is worth noting an alternative way to express $c^L(o)$. For each confidence level o, 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, using the vector notation introduced in Example 3.2, as:

$$c^L(o) = \left\{ \mathbf{q} \in [0,1]^3 : (\mathbf{q} - \mathbf{p^L})^T \mathbf{D^L} (\mathbf{q} - \mathbf{p^L}) \leqslant o \right\} \tag{7}$$

where

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

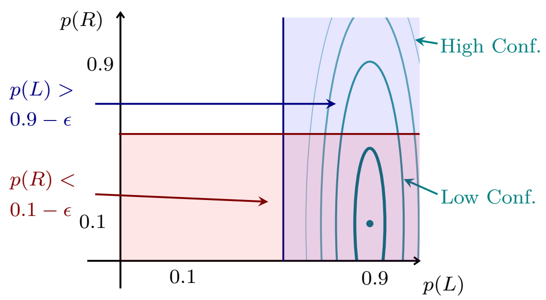

(a) Illustration of Proposition 1.
Note: As in Figure 2a, the blue area represents the probability judgement that $p(L)$ is within ϵ of Laura's best-guess probability $p^L(L) = 0.9$; the red area represents the judgement that $p(R)$ is within ϵ of $p^L(R) = 0.1$. The confidence in these judgements corresponds to the largest elliptical set contained in each area (Section 2.1): it is higher for the judgement concerning L.

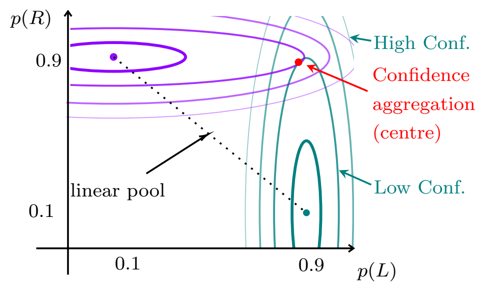

(b) Illustration of expert-sensitive pooling (Example 3.4).
Note: The red point is the centre of the result of confidence aggregation applied to the two confidence rankings, which coincides with expert-sensitive pooling (Definition 4). The aggregate probability of L is closer to Laura's judgement ($p^L(L) = 0.9$), 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 $w_L^L > w_R^L$ and w_B^L 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 $w_L^R < w_R^R$ and w_B^R low.

and similarly for Ray, with

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

3.3 Aggregation with within-person expertise diversity

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

Theorem 2. *Suppose that each agent $i = 1, \ldots, n$ has a confidence ranking of the form (5), with 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:*

$$\arg\min_{p \in \Delta} \sum_{i=1}^{n} \sum_{j=1}^{m} w_j^i d(p|_{\mathcal{P}_j}, p^i|_{\mathcal{P}_j}) \tag{8}$$

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

$$\iota(q) = \sum_{j=1}^{m} w_j d(q|_{\mathcal{P}_j}, p|_{\mathcal{P}_j}) \tag{9}$$

We shall examine the properties of this aggregate judgement presently. Before doing so, we note that confidence aggregation applied to $\mathbf{w}$ d-confidence rankings generates a new family of pooling rules.

To this end, let us define the function yielding the lowest-confidence set 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 distance. The function $F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m} : \Delta^n \to 2^\Delta$ is defined by

$$F^d_{\mathcal{P}_1,\ldots,\mathcal{P}_m}(p^1, \ldots, p^n) = \arg\min_{p \in \Delta} \sum_{i=1}^{n} \sum_{j=1}^{m} w_j^i d(p|_{\mathcal{P}_j}, p^i|_{\mathcal{P}_j}) \tag{10}$$

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—a rule taking a profile of probability measures $(p^1, \ldots, p^n) \in \Delta$ to a probability measure. 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, note that, given the previous definitions, the centre of the aggregate confidence ranking (8) can equivalently be characterised as the set of probability measures p such that $(p|_{\mathcal{P}_1}, \ldots, p|_{\mathcal{P}_m})$ belongs to:

$$\arg\min_{(p_1,\ldots,p_m) \in P_{\mathcal{P}_1,\ldots,\mathcal{P}_m}} \sum_{i=1}^{n} \sum_{j=1}^{m} w_j^i d(p_j, p^i|_{\mathcal{P}_j}) \tag{11}$$

Whenever d is convex, (11) 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, (8) 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. Hence we have the following Proposition.

Proposition 2. *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.*

Hence confidence aggregation generates this new well-defined pooling rule, which we call *expert-sensitive pooling*. As we now show on an example, it incorporates within-person cross-issue expertise diversity in a natural way.

Example 3.4. Suppose that Laura and Ray have the confidence rankings defined in Example 3.3 with $w_L^L > w_R^L$ and $w_L^R > w_R^R$. 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. Since the example stipulates that Laura has more expertise in the labour market than Ray, it is natural, in the light of the analysis of confidence rankings of form (5), to assume that $w_L^L > w_L^R$. Similarly, given Ray's higher specialisation in the real estate sector, $w_R^R > w_R^L$.

Using the formulation of the image of these confidence rankings in $\Delta(\mathcal{P}_L) \times \Delta(\mathcal{P}_R) \times \Delta(\mathcal{P}_B)$ (Remark 1), and integrating the constraints defining $P_{\mathcal{P}_L,\mathcal{P}_R,\mathcal{P}_B}$, as specified in Example 3.2, the minimisation problem (11) defining the centre of the confidence aggregation becomes:

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

where $\mathbf{A}$ and $\mathbf{r}$ are as defined above (Example 3.2).

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:

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

where, as specified above, $p^L(L) = p^R(R) = 0.9$, $p^L(R) = p^R(L) = 0.1$ and $p^L(B) = p^R(B) = 0.09$.

Note that $\frac{w_B^L+w_B^R}{w_L^L+w_L^R}$ 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 $\frac{w_B^L+w_B^R}{w_R^L+w_R^R}$. 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 group probabilities tend to:

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

So the group 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 $w_L^L > w_L^R$, the group 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 agents' 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 group 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 low on either ranking, and this picks out probability measures that are close to both Laura's probability on L and Ray's on R. Hence, as desired, confidence aggregation applied to these confidence rankings, which properly reflect cross-issue expertise differences, yields a group judgement that follows each agent more closely on her area of expertise, as one would have wanted. Moreover, this also shows that expertise-sensitive pooling (Definition 4) can reflect these within-person cross-issue expertise differences, and hence fairs better on this score than linear (or, for that matter, geometric) pooling.

Given the confidence rankings of the form (6), where the centre is a probability measure 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 $p(B) = 0.09$ (i.e. when the constraints (4) are slack), 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 which satisfies the constraints, i.e. $p(L) + p(R) - 1$. Since this is typically not 0.09[9], this example demonstrates that the confidence aggregation rule does not respect spurious unanimities. The same goes for expertise-sensitive pooling.

[9]E.g. when $w_L^L = w_R^R = 0.75$, $w_R^L = w_R^L = 0.25$, $p(L) = p(R) = 0.7$, and $p(B) = 0.4$.

The case where $\frac{w_B^L+w_B^R}{w_L^L+w_L^R} \to 0$ and $\frac{w_B^L+w_B^R}{w_R^L+w_R^R} \to 0$ translates low confidence in the judgements about B compared to the judgements concerning L and R. This is clearly most relevant to the example in the Introduction, where Laura's expertise, say, concerns L but not R, so there is no reason to expect her to have particular expertise on $B = L \cap R$. To complete the discussion, 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:

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

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.

This example shows that, when there is comparative expertise on issues, both confidence aggregation and expertise-sensitive pooling faithfully reflect it in the resulting group probability judgements. As such, they resolve the within-person expertise diversity challenge. Moreover, in so doing, the example also shows that these 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.

The early literature on probabilistic belief aggregation contains suggestions using weighted averaging with potentially different weights for each event (e.g. Bordley and Wolff, 1981). To a certain extent, the limit case in the example above captures the intuition behind these proposals, for the events L and R. More importantly, it overcomes their well-understood limits. Such rules are not 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). However, the expertise-sensitive pooling rule derived from confidence aggregation is well-defined, by Proposition 2; accordingly, it does not coincide with weighted averaging for all events. The event B is a clear example of this: when $p(B) = p(L) + p(R) - 1 > 0.09$, it is not a weighted average of the individuals' probability judgements for B.

Example 3.4 illustrates a straightforward application of confidence aggregation, with rankings reflecting varying expertise, to the example in the Introduction. One central factor is the trade-offs between the confidence in judgements concerning the main two issues—labour and real estate—and what happens to both, considered as a third issue. However, an alternative possibility for analysing this example is to consider that the individuals have opinions on the main issues and their relationship, rather than 'primitive' views on B. We now show that the confidence approach can easily cope with such possibilities.

Example 3.5. Here 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,[10] without being maximally confident in this judgement. Note that the belief in independence implies that $p^L(B) = p^R(B) = 0.09$, as per Table 1. To integrate this, take a 3-dimensional 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 consider the following confidence ranking:

$$c^L_{Ind}(o) = \left\{ q \in \Delta : \begin{array}{c} \sum_{j=\{L,R\}} w^L_j (q(j) - p^L(j))^2 \\ + w^L_I \left(q(B) - q(L).q(R)\right)^2 \end{array} \leqslant o \right\} \tag{12}$$

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

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

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

Here the aggregation on each of the issues L and R uses issue-specific weights, reflecting differing confidence, as in the limit case in Example 3.4. For the issue B, the assumption that agents consider independence to be their best guess concerning the relationship between L and R generates the probability.

[10] In the standard probabilistic sense: $p^i(L \cap R) = p^i(L)p^i(R)$.

In tandem with the previous example, 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.[11] This is relevant for another recurrent criticism of linear pooling: that is 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 leading example: the linear pool of Laura's and Ray's probabilities with equal weights ($w^L = 1-w^L = \frac{1}{2}$) 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 (12) shows how confidence aggregation can respect independence, whilst retaining much of the spirit of linear pooling. For instance, when $w^L_L = w^L_R = w^R_L = w^R_R$, the resulting centre probability is $p^{LP}(L) = 0.5$, $p^{LP}(R) = 0.5$, $p^{LP}(B) = 0.25$: i.e. the same as linear pooling for the issues L and R, but with independence retained (and hence a different B).

The beliefs about the independence of L and R in Example 3.5 are considered merely for the purposes of illustration.[12] The point of the example is more general: by incorporating conditional probabilities in much the way proposed in Eq. (12), 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 confidence aggregation. We begin by setting out the framework and the representation of preferences.

4.1 Preferences

We use a standard Anscombe-Aumann-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, as in the Anscombe and Aumann (1963) setting. $\mathcal{A}$ is the set of *acts*: (measurable) functions from states Ω to consequences $\mathcal{X}$. $\mathcal{A}^c$ is the set of constant acts (acts taking a constant value). Mixtures of acts are defined

[11] 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.

[12] Arguably, the effects of an interest rise on the labour and real estate sectors would not typically be considered independent.

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 by $f_\alpha g$, 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 $\succ$ will be said to *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 discussed in Hill (2019b), there are several decision models in the confidence family. Here we work with the incomplete preference version of the confidence model (Hill, 2016), with strict preferences as primitive, as in Bewley (1986); Galaabaatar and Karni (2013).[13] According to it, for all acts $f, g \in \mathcal{A}$, $f \succ g$ if and only if:

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

where $\mathbb{E}_p$ is the expectation with respect to a probability measure $p \in \Delta$,[14] $u : \mathcal{X} \to \mathbb{R}$ is a non-constant affine utility function, c is a closed confidence ranking and D is a function from $\mathcal{A} \times \mathcal{A}$ to O, satisfying the following *richness* condition: for every $(f,g) \in \mathcal{A} \times \mathcal{A}$ and $o \in D(\mathcal{A} \times \mathcal{A})$, there exist $h \in \mathcal{A}$ and $\alpha \in (0,1]$ such that $D(f_\alpha h, g_\alpha h) = o$. This function, called the *cautiousness coefficient*, picks out the relevant confidence level for the decision, and captures the decision maker's ambiguity attitudes, or attitudes to choosing on the basis of limited confidence. We refer to the cited papers for discussion and details.

When (13) holds for preferences $\succ$, we say that the triple (c, D, u) *represents* $\succ$. In this case, there is a unique u, up to positive affine transformation, a unique closed and convex c and a unique minimal D representing $\succ$ (Hill, 2016); we refer to this (c, D, u) as *the representation of* $\succ$.

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. We assume that all preferences are represented according the confidence model (13).[15]

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

Behavioural foundations for an incomplete preference version of the confidence model have been provided in Hill (2016).[16] They can be used to provide a reformulation of this assumption in terms of preferences.

[13] Given the close relationship between incomplete and ambiguity averse preferences (Ghirardato et al., 2004; Gilboa et al., 2010), similar foundations to those developed here can be provided in terms of the maxmin-EU member of the confidence family (Hill, 2013), for instance.

[14] That is, for any $\phi : \Omega \to \mathbb{R}$, $\mathbb{E}_p \phi = \sum_{\omega \in \Omega} p(\omega)\phi(\omega)$.

[15] Note that this implies that preferences are non-contradictory.

[16] The cited paper takes the weak preference relation as primitive; similar techniques can be used with the strict preference, drawing on the work of Bewley (1986); Karni (2011).

Following Assumption 1, let (c^0, D, u) be the representation of $\succ^0$, and, for each $i = 1, \ldots, n$, (c^i, D^i, u^i) be the representation of $\succ^i$. To focus on aggregation of beliefs, we follow other papers on that topic (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 and the cautiousness coefficient—this is expressed by the following assumption.

Assumption 2. *For all $i = 1, \ldots, n$:*

1. *u^i and u are identical up to positive affine transformation;*
2. *$D^i = D$.*

A central idea behind the confidence model 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, 2019b). In the light of this, when higher-confidence beliefs are invoked—i.e. $D(f,g) > D(f',g')$—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. 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. We use $s, s' \in \mathcal{S}$ as notation for stakes levels. Hill (2016) contains several examples of such (real-valued) notions of stakes. For $(f,g) \in \mathcal{A} \times \mathcal{A}$ and $s \in \mathcal{S}$, we say that (f,g) *has stakes* s (read as: 'the choice between f and g has stakes s') if $\sigma(f,g) = s$. Under the intuition mooted above, the confidence level deemed appropriate for a decision is determined by the stakes involved in the decision: formally, this amounts to the assumption that there exists a monotonically increasing function $\zeta : \mathcal{S} \to O$ such that $D = \zeta \circ \sigma$. In general, different agents could use the same notions of stakes σ but have different cautiousness coefficients D^i, corresponding to different ζ^i. However, in the context of Assumption 2, the cautiousness coefficient is the same for all agents; for simplicity, we thus assume that ζ is the identity, so $D = \sigma$.

Finally, given a preference relation $\succ$ represented according to (13) and a stakes level $s \in \mathcal{S}$, we define the derived relation $\succ_s$ 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 $f_\alpha h \succ g_\alpha h$.[17] 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

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

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 indetermine 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, then f would typically be chosen over g: i.e. $f \succ_s g$, where s is the appropriate stakes level. 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.

Whilst stated on the models for ease, Assumption 2 can be reformulated in behavioural terms. The first clause corresponds to the standard axiom that all individuals and the group have the same preferences over constant acts. The latter is built into axiomatisations of the confidence model assuming an exogenously given notion of stakes (Hill, 2016); framework-specific axioms characterise it in setups where stakes are endogenous (Hill, 2015).

We shall use vector notation, and often express the tuple of stakes levels $(s_1, \ldots, s_n) \in \mathcal{S}^n$ as the vector $\mathbf{s}$. Under this notation, s_i is understood to be the i^{th} stakes level under vector $\mathbf{s}$. The following definition shall play a central role in the sequel.

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 indivuduals in the group, at the specified stakes levels. The group exhibits consensus across the tuple of stakes levels $\mathbf{s}$ if the assembled preferences are consistent; 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 $\succ_s^0$ appear in the consensus, though some preferences determined in the consensus may be left open in $\succ_s^0$. In other words, consensus respect 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 (Strict (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 at 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 Strict (issue-wise) Pareto nor Corpus-wise Pareto imply the other, Theorem 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.*

Under 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 be 'inherited' from a level where some are. Non-degeneracy states that there is some consensus among the individuals: if they leave sufficiently many preferences aside, they can come to a consensus.

Our base characterisation result is that the previous axioms characterise consensus-preserving confidence aggregation.

Theorem 3. *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 confidence-level aggregator $\otimes$ *that is minimal on consensuses 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, even Consensus-based beliefs can be dropped without jeopardising the core of confidence aggregation: in its absence, the group confidence ranking is that obtained by a confidence aggregation, except at confidence levels at the bottom of the ranking.

As indicated previously, 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 confidence-level aggregator: 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 corresponding to the confidence-level aggregators mentioned in Section 2.2; details are given in Appendix A.

5 The Performance of Confidence Aggregation

A wide body of literature in cognitive psychology has studied people's confidence in their answers or actions in a wide variety of situations, including those involving perception, memory, cognitive tasks and decision making (Metcalfe and Shimamura, 1994; Koriat, 2012a; Rahnev et al., 2020). Some empirical evidence in that literature even suggests the communicating confidence judgements—as broadly recommended by our proposed aggregation procedure—can improve group decisions (Bahrami et al., 2010), at least in some situations (Koriat, 2012b). In this section, we draw upon well-known findings in this literature to propose a rudimentary comparison of the *performance* of confidence aggregation against that of linear pooling—where performance is measured by the propensity to yield group judgements that are closer to the truth.

We work with a special case of the setup from Section 3.2, with two-element partitions. There are n individuals, each of whom provides a probability judgement on each of m issues, where an issue is an statement that is true (1) or false (0). For ease, we assume that the issues are logically independent, in the sense that any combination of truth values for each issue is possible. So each individual's report can be any point in $[0,1]^m$, and the group must come to a (probabilistic) judgement on each of the issues, which too will correspond to a point in $[0,1]^m$. Performance of an aggregation rule will be assessed by the result's proximity to the truth. We adopt the convention that for each issue, the statement corresponding to it is true, so the closer a judgement is to 1, the better it is, for all issues. The truth is thus $\mathbf{1}$, the vector of 1's.

Any evaluation of the performance of aggregation rules will have to rely on some understanding of where individuals' confidence assessments come from, or at the least how they are related to the correctness of the stated probability judgements. Here we rely on insights from the psychology literature, which typically focusses on subjects' confidence concerning their performance in a task. For instance, the task could be to answer a cognitive or knowledge question, and the subject would then report her confidence that his reply was correct, which may or may not be a probability (e.g. Rahnev et al., 2020). The 'task' of each expert in a standard probability aggregation situation is to state informed subjective probabilities for the issues at hand—this is the data used by pooling rules. In the light of this relationship, the comparison with confidence aggregation amounts to seeing what changes when experts also report their confidence in their probability judgements.

The psychology and neuroscience literature has produced an impressive array of (competing) models of the sources of confidence judgements. Most build upon existing models of task responses, which are generally stochastic (e.g. Green and Swets, 1966; Ratcliff, 1978; Mamassian, 2016). Although, to our knowledge, there have been no studies of confidence concerning probability reports in the psychological literature—and hence models which could apply directly here[18]—we work with the following simple but fairly general model of the cognitive process 'behind' confidence and probability reports.

Reflecting the stochastic nature of the cognitive processes underlying reports, we assume that each individual i draws a vector $\mathbf{p^i}$ of probability values or 'signals' for each of the issues from a distribution over $[0,1]^m$, with mean $\mu^{\mathbf{i}}$ and positive-definite covariance matrix $\mathbf{\Gamma^i}$. This distribution captures the accuracy of and noise in the cognitive process of the individual—where we understand the notion of 'cognitive process' in this context to be wide enough to encompass the (perhaps differing) information held by the individuals. The individual reports $\mathbf{p^i}_j$ as her best-guess probability for each issue j. We assume moreover than the draws are independent both across individuals and across issues.

Assumption 3. *For all $i, i' \in \{1, \ldots, n\}$ and $j, j' \in \{1, \ldots, m\}$ with $i \neq i'$ or $j \neq j'$, $\mathbf{p^i}_j$ and $\mathbf{p^{i'}}_{j'}$ are independent.*[19]

This translates the assumption that the stochastic cognitive processes determining the probability reports individuals bring to the table are independent across individuals

[18] Most models, to our knowledge, concern tasks with few alternatives, often two, while the task of reporting a probability can in principle yield any real number between 0 and 1.

[19] So, in particular, $\mathbb{E}(\mathbf{p^i}_j \mathbf{p^{i'}}_{j'}) = \mathbb{E}(\mathbf{p^i}_j)\mathbb{E}(\mathbf{p^{i'}}_{j'})$.

and issues. It allows for interactions in beliefs across individuals and issues, as long as they translate into relations between properties of the distributions governing the stochastic processes, rather than into dependencies in the processes themselves. For instance, if Ann considers the belief about one issue to be related to that about another, this could be reflected by appropriate relationships between, say, the means of her distribution for the two issues. Or if Bob's beliefs about an issue are influenced by Cat's, this can translate into a specific link between the means or variances of their distributions on the given issue. Under Assumption 3, such relations cover all relevant interaction between beliefs: there is no further correlation between the 'signals' drawn at the moment of probability reporting. In Appendix B.3, we prove a stronger version of the main result in this section, where this independence assumption is weakened.

In the light of Assumption 3, each $\mathbf{\Gamma^i}$ is a diagonal matrix, with the variances of signals for each issue, $(\sigma^i_j)^2$, on the diagonal.

Turning to confidence, each individual's confidence is represented by a confidence ranking centred on her reported best-guess probability $\mathbf{p^i}$. We assume that the confidence ranking is generated by the Mahalanobis distance from $\mathbf{p^i}$, according to the inverse of a positive-definite matrix $\mathbf{\Sigma^i}$. We also assume that $\mathbf{\Sigma^i}$ is diagonal for each i, with entries $(\rho^i_j)^2$ on the diagonal. In other words, the confidence ranking corresponds to the implausibility function given by[20]

$$\iota(\mathbf{x}) = (\mathbf{x} - \mathbf{p^i})^T \mathbf{\Sigma^i}^{-1} (\mathbf{x} - \mathbf{p^i}) \tag{14}$$

$\mathbf{\Sigma^i}$ may coincide with the covariance matrix of the process determining probability reports $\mathbf{\Gamma^i}$, though this need not be the case. Each individual i reports her best-guess probability $\mathbf{p^i}$ and her confidence ranking, determined by $\mathbf{p^i}$ and $\mathbf{\Sigma^i}$, so $\big((\mathbf{p^1}, \mathbf{\Sigma^1}), \ldots, (\mathbf{p^n}, \mathbf{\Sigma^n})\big)$ characterises the profile of probability-and-confidence reports across all individuals.

Our main substantive assumption draws on two insights from the aforementioned psychology literature. The first is the close, inverse relationship between the confidence in one's response to a task and the variability of that response. This is supported by empirical studies (Koriat, 2012a), and many models of confidence judgements relate, in some way or another, confidence judgements to the signal(s) or stochastic process which determines the answer or action in the task, either because they monitor actual signal production, or a continuation of the signal production mechanism, or are determined by the signal structure (e.g. Pleskac and Busemeyer, 2010; Ratcliff and Starns, 2013; Mamassian, 2016). The second is the old and well-established finding that in many situations most people have 'positive metcognitive sensitivity' (Henmon, 1911; Fleming and Lau, 2014; Rahnev et al., 2020) or 'monitoring resolution' (Koriat, 2012a): the

[20] Note that this belongs to the family of expertise-sensitive confidence rankings defined in Section 3.2; see Definition 3 and Remark 1.

confidence they report is positively correlated with their performance in the task. In the case of interest here, this means that an individual whose probability judgements on one issue are typically closer to the truth that those on another issue—she tends to perform better in the task concerning the first issue—will tend to have higher confidence in the former judgements. These inspire the following assumption.

Assumption 4. *For all i, i', j, $(\sigma^i_j)^2 \geqslant (\sigma^{i'}_j)^2$ if and only if $(\rho^i_j)^2 \geqslant (\rho^{i'}_j)^2$ if and only if $1 - \mu^{\mathbf{i}}{}_j \geqslant 1 - \mu^{\mathbf{i}'}{}_j$.*

The assumption that individuals whose probability reports are on average closer to the truth (which, recall, is 1) on an issue j also have lower $(\rho^i_j)^2$ on that issue—and hence more confidence in their judgements about it—reflects the finding that confidence assessments accurately distinguish between correct and incorrect judgements. The assumption that $(\rho^i_j)^2$ is lower—i.e. confidence is higher—when the variance $(\sigma^i_j)^2$ in the signal underlying the report is lower reflects the insight behind the neuro-psychology models than confidence is inversely related to the dispersion produced by the mechanism yielding the task response.

The finding of positive metacognitive sensitivity is not to be confused with the phenomenon of 'metacognitive bias' or 'monitoring miscalibration' (Fleming and Lau, 2014; Koriat, 2012a): the 'absolute' confidence levels may be misaligned with the actual performance in the task. The well-known overconfidence bias is an example of this: for instance, when a subject states as confidence a probability of success in the task which is larger than her actual success rate.[21] People can be overconfident—expressing higher than warranted confidence in performance—whilst also exhibiting positive metcognitive sensitivity—they express higher confidence in tasks which they perform better in. Assumption 4 may still hold if individuals have metacognitive bias, and even if different individuals exhibit different amounts of metacognitive bias. It only requires that the metacognitive bias does not outweigh the positive metcognitive sensitivity: if Ann is more of an expert than Bob on issue 1, and B is more of an expert on issue 2, then metacognitive sensitivity implies that Ann is more confident in her beliefs about 1 than those about 2, and vice versa for Bob. Our assumption requires that, no matter how underconfident she in general and how overconfident Bob is, she remains more confident in her beliefs about 1 than Bob is in his beliefs about 1. We take it that the appropriateness of this assumption may depend on the situation, and read the result below as drawing out consequences for performance in situations where the assumption holds,

[21]Note that such biases can typically only be measured if confidence is elicited on a probability scale. This scale is not used in many psychology studies (Fleming and Lau, 2014); moreover, confidence rankings, since they are ordinal (Section 2; Hill, 2013) do not necessarily permit measurement on a probability scale. So standard notions of overconfidence do not apply, as such, to them.

without meaning to claim that all situations must be of this sort. We leave a systematic exploration of the circumstances in which such biases are more or less propitious for confidence aggregation for future research.

We compare two pooling rules. The first is (equal-weight) linear pooling, defined by:

$$\lambda((\mathbf{p^1}, \mathbf{\Sigma^1}), \ldots, (\mathbf{p^n}, \mathbf{\Sigma^n})) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{p^i} \tag{15}$$

Note that no assumption of general expertise over all issues is made here—indeed, each individual could be an expert on a different issue. Moreover, it is not assumed that the group (or its representative) has any information on the relative expertise of members—such information could trump reported confidence, hence requiring a more refined analysis. There thus seems to be no reason to use anything other than the equal-weight version of linear pooling for comparison. Section 7.3 discusses some extensions of the result obtained here to weighted linear pooling.

The other rule is confidence aggregation under the average confidence-level aggregator, with the centre of the aggregate confidence ranking used as the group judgement, as set out in Section 3 (see Figure 1). Formally:[22]

$$\phi((\mathbf{p^1}, \mathbf{\Sigma^1}), \ldots, (\mathbf{p^n}, \mathbf{\Sigma^n})) = \arg\min_{\mathbf{x}} \sum_{i=1}^{n} (\mathbf{x} - \mathbf{p}^i)^T \mathbf{\Sigma^i}^{-1} (\mathbf{x} - \mathbf{p}^i) \tag{16}$$

We will compare these two pooling rules using two performance measures. For a vector of probability judgements, one for each issue, $\mathbf{p}$, its absolute distance from the truth on issue j is $1 - \mathbf{p}_j$ if issue j is true, and $\mathbf{p}_j$ otherwise. So, in our setup where all issues are true, the *mean absolute distance* of a judgement $\mathbf{p}$ from the truth, taken over all issues, is $\sum_{j=1}^{m}(\mathbf{1} - \mathbf{p})_j$.

On the other hand, a popular measure used for evaluating expert assessments in theory and in practice is the *Brier score* (Brier, 1950; Winkler and Murphy, 1968; Cooke, 1991). In our setup, it is defined, for a vector of probability judgements $\mathbf{p}$, as $B(\mathbf{p}) = \sum_{j=1}^{m}(\mathbf{1} - \mathbf{p})_j^2$. The lower the Brier score, the lower the Euclidean distance to the truth.

Our main result says that, in expectation, confidence aggregation yields a group judgement that is closer to the truth than linear pooling, under both of these measures.

Theorem 4. *Under Assumptions 3 and 4:*

i in expectation, the mean absolute distance from the truth is smaller under confidence aggregation with the average confidence-level aggregator, as compared to linear pooling:

$$\mathbb{E}\sum_{j=1}^{m}\big(\mathbf{1} - \phi((\mathbf{p^1}, \mathbf{\Sigma^1}), \ldots, (\mathbf{p^n}, \mathbf{\Sigma^n}))\big)_j \leqslant \mathbb{E}\sum_{j=1}^{m}\big(\mathbf{1} - \lambda((\mathbf{p^1}, \mathbf{\Sigma^1}), \ldots, (\mathbf{p^n}, \mathbf{\Sigma^n}))\big)_j$$

[22] By Proposition 2, this is a well-defined probability aggregation rule.

ii in expectation, the Brier score is smaller under confidence aggregation with the average confidence-level aggregator, as compared to linear pooling:

$$\mathbb{E}B(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))) \leqslant \mathbb{E}B(\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))$$

This Theorem suggests that, under an assumption that translates findings in psychology about the role of confidence assessments and their relationship to probability reports, confidence aggregation outperforms linear pooling: in expectation, it provides judgements that are closer to the truth. So, not only does it overcome the challenges for linear pooling discussed in previous sections, it promises to fair better epistemically, at least in a range of situations. The basic intuition behind this result is not completely disassociated with how confidence aggregation incorporates within-person expertise diversity, as discussed in Section 3. Since confidence in a probability judgement co-varies with its accuracy, by giving judgements on issues where individuals have higher confidence more weight, confidence aggregation tends to produce judgements that are closer to the truth.

6 Confidence aggregation and dynamic rationality

This section treats a common theme in the aggregation literature, which has recently been reintroduced by Dietrich (2021): the interaction between aggregation and update. Dietrich argues that a 'rational group' requires belief aggregation to be in sync with belief updating. This is typically formulated in terms of commutation 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. And Bayesian conditionalisation no longer applies, without revision, to such beliefs. Hill (2022) proposes a *confidence update* rule for confidence in beliefs, and argues for its normative validity, suggesting in particular that it deals appropriately with situations in which standard Bayesian update struggles. So the question of dynamic rationality in our context is whether confidence aggregation, as set out in Section 2.2, commutes with confidence update.

In the framework set out in Section 2, 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, $\mathcal{C}_E = \{p(\bullet|E) : p \in \mathcal{C}, p(E) > 0\}$:

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

$$\Phi(o) = \{p \in c(o) : p(E) \geqslant \rho_E(o)\}_E \tag{17}$$

Readers are referred to 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 5. *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:*

$$F_{\otimes}(c_1|\rho_E, \ldots, c_1|\rho_E) = F_{\otimes}(c_1, \ldots, c_n)|\rho_E \tag{18}$$

So confidence aggregation commutes with confidence update: it is 'dynamically rational' with respect to the appropriate update rule for confidence, to use the term coined by Dietrich (2021). As argued by Dietrich and others, such coherence can be considered 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. This Theorem thus provides a reassuring message concerning confidence aggregation's credentials on this score.

7 Discussion

7.1 Probabilistic and non-probabilistic belief aggregation

The developments in this paper have largely focused on the contribution of confidence aggregation with respect to recognised challenges for aggregating probability measures. Part of the literature on probability aggregation takes probabilities as primitive, rather than working with (subjective expected utility) preferences; a classic survey is Genest and Zidek (1986). The within-person expertise diversity challenge has been raised in this literature, as has already been discussed in Section 3.

The confidence aggregation rule defined in Section 2.2 operates directly on confidence rankings, and hence, like pooling rules, does not require a preference setup to be applied. However, just as pooling rules tacitly assume interpersonal comparison of probability judgements—i.e. one can say when two agents are assigning the same probability—in direct application, confidence aggregation assumes interpersonal comparison of confidence levels: i.e. one can tell when two agents are talking about the same confidence level. (In the preference foundations provided in Section 4, this is implied by the assumption that all agents have identical tastes.) Hill (2019a) provides a detailed discussion of the problem of 'calibrating' confidence levels across individuals, and provides a calibration scale, drawing on an analogy with the role the probability scale plays in calibrating subjective probability judgements. Such a scale can be used for applications of confidence aggregation.

Another part of the literature touching 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 also for aggregation of belief *tout court*. For instance, Mongin and Pivato (2020); Dietrich (2021) 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 the 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 6. Mongin and Pivato (2020) restrict Pareto to objective uncertainty; indeed, their representation involves 'no connection between the social probability and the individual ones'. Unlike the approach developed here, these make no attempt to retain the insight behind Pareto by replacing it with a more appropriate consensus preservation condition. By contrast, Bommier et al. (2021) present a condition involving the preservation of consensuses concerning prospects yielding identical distributions of outcomes for all individuals, and use it to provide a decision rule aggregating probabilistic individual 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. Drawing on Gilboa et al. (2004), Alon and Gayer (2016) consider aggregation of subjective expected utility preferences, where group preferences may be non-expected utility; Stanca (2021) undertakes a similar exercise, for a different class of non-expected utility preferences, and under the assumption that all individuals have the same utilities. Both involve versions of Pareto that, were group preferences expected utility, would lead to linear pooling.

To the extent that confidence rankings support both classes of ambiguity averse and incomplete preferences (Hill, 2013, 2016), 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, whereas Nascimento (2012); Hill (2012) characterise similar rules for more general classes of preferences, all under the assumption of identical utilities. Danan et al. (2016) explore aggregation of incomplete preferences, with potentially differing utilities and beliefs. All these approaches 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). It 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.[23]

As recalled in Section 4.1 (see also Hill, 2013, 2019b), the set of priors in the confidence ranking that are mobilised in a decision will depend on the stakes involved, with larger sets being used for more important decisions. So, whilst the group may act as a subjective expected utility agent with the aggregate (centre) probability when the stakes are low, when the stakes are higher, it will use a larger set of priors, and, in the case of the maxmin-EU version of the model (Hill, 2013), be more ambiguity averse. Importantly, the aggregation procedure is independent of the stakes relevant for any post-aggregation decision, so the same aggregate confidence ranking can be used for both low- and high-stakes decisions. One domain where such flexibility could be useful is where the 'experts' are 'models'.

7.2 Aggregating models: averaging and misspecification

In a range of domains, including climate science and economics, decision makers are faced with a set of (scientific) models, each of which may give different predictions or evaluations of prospects. One popular approach in such contexts is Bayesian Model Averaging (Raftery et al., 1997; Steel, 2020). At its base, it involves linear pooling of the probability distributions provided by the various models, with weights determined by the posterior probabilities over the models. As such, Theorem 1 shows that Bayesian Model Averaging is a special case of confidence aggregation, corresponding to particular assumptions about the confidence in the various models. Moreover, the subsequent developments in Section 3, notably the development of probability aggregation rules accommodating within-person expertise diversity, may be relevant when the 'experts' are 'models'. In climate science for instance, it is not uncommon for some models to be 'better' on certain issues, and others 'better' on others. One model could have a more detailed representation of cloud formation, whereas another is more accurate on elements

[23]This can be seen from the fact that Nau's rule violates (3); see Nau (2002, Figs 2 & 3).

of the biosphere: the former might thus be expected to do a better job in predicting hurricanes, and the latter in predicting ground-level temperature. The developments in Section 3.3 provide a blueprint for procedures for aggregating models that can faithfully integrate inherent differences in models' domains of specialisation.

Another direction taken in the face of a class of models focuses on robustness; in the economics literature, this is typically associated with decision rules sensitive to the possibility of model misspecification (Hansen and Sargent, 2001; Hansen, 2007; Hansen and Sargent, 2022). As ambiguity averse decision models, they do not coincide with subjective expected utility with a unique probability measure—which would be the natural decision procedure associated with model averaging. The confidence approach can nevertheless recoup the essence of this misspecification approach too, and sheds light on its relationship to model averaging. We illustrate this point on an example.

Consider a set $\mathcal{M} \subset \Delta$ of models, each of which is a probability distribution (over states). Let $w : \mathcal{M} \to \Re_{\geqslant 0}$ be an assignment of weights to models, and consider the w reverse relative entropy confidence rankings generated by each model (Definition 2): for $m \in \mathcal{M}$, $c_m(o) = \{q \in \Delta : w(m)R(m\|q) \leqslant o\}$. Under confidence aggregation with the average confidence-level aggregator (Example 2.2), the group confidence ranking $c_{\otimes^{av}}$ is such that, for every confidence level o, $c_{\otimes^{av}}(o) = \{q \in \Delta : \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} w(m)R(m\|q) \leqslant o\}$, where in this example we take $\mathcal{M}$ to be finite for ease.[24] Under the maxmin-EU version of the confidence model (Hill, 2013), an act f is thus evaluated according to

$$\min_{q \in \Delta : \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} w(m)R(m\|q) \leqslant D(f)} \mathbb{E}_q u(f)$$

where $D : \mathcal{A} \to O$ is a cautiousness coefficient assigning the appropriate confidence level for evaluating act f on the basis of the stakes involved. By Theorem 1, for low enough stakes, this becomes

$$\mathbb{E}_{\sum_{m \in \mathcal{M}} \frac{w(m)}{\sum_{m \in \mathcal{M}} w(m)} m} u(f)$$

i.e. the decision maker evaluates acts according to the model average, with weights $\frac{w(m)}{\sum_{m \in \mathcal{M}} w(m)}$. For appropriate w, this coincides with the evaluation of acts, and hence choice, under Bayesian Model Averaging. On the other hand, for higher stakes, say with $D(f) = o$, an act f is evaluated according to

$$\min_{q \in \Delta : \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} w(m)R(m\|q) \leqslant o} \mathbb{E}_q u(f)$$

where $\{q \in \Delta : \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} w(m)R(m\|q) \leqslant o\}$ is not a singleton. This translates greater aversion to ambiguity, or concern for misspecification of the models. For medium stakes,

[24] Extensions to the infinite case can be carried out with known techniques.

the set may not be too far away from the Bayesian model average; for very high stakes, and hence high o, the set over which the minimisation is taken may include $\mathcal{M}$. As might be expected, more robustness to misspecification is involved when the decision is more important.

A tighter connection to existing misspecification models comes when combining the same maxmin-EU version of the confidence decision model with the 1 relative entropy confidence rankings (Definition 2)—i.e. $c_m(o) = \{q \in \Delta : R(q\|m) \leqslant o\}$ for each $m \in \mathcal{M}$—and the *minimum aggregator*. The latter is defined by $\otimes^{min}\mathbf{o} = \min_{m\in\mathcal{M}} o(m)$. This yields a decision maker who evaluates acts according to:

$$\min_{q\in\Delta:\min_{m\in\mathcal{M}} R(q\|m)\leqslant D(f)} \mathbb{E}_q u(f) = \min_{m\in\mathcal{M}} \left(\min_{q\in\Delta:R(q\|m)\leqslant D(f)} \mathbb{E}_q u(f) \right) \tag{19}$$

For a fixed $D(f)$ and a singleton $\mathcal{M}$, these are just the constraint preferences defined by Hansen and Sargent (2001, 2008). So this application of confidence aggregation in the context of the confidence decision model extends constraint preferences, firstly, by centring on a set of models rather than a single one, and secondly, by allowing the degree of concern for misspecification—$D(f)$ in the expression above—to depend on the importance of the decision. The second aspect has already been discussed in Hill (2019b, Sect 5); here we focus on the first.

Hansen and Sargent (2008) show that constraint preferences yield the same optimal behaviour as so-called multiplier preferences on various classes of decision problems. Whenever $\mathcal{M}$ is a convex, compact set, essentially the same proof can be used to show that, for these classes of decision problems, the optimal choice under (19) for fixed $D(f)$ coincides with the optimal choice under:[25]

$$\min_{q\in\Delta} \left(\mathbb{E}_q u(f) + \lambda \min_{m\in\mathcal{M}} R(q\|m) \right) \tag{20}$$

for appropriate λ. Multiplier preferences (Hansen and Sargent, 2001) correspond to the special case where $\mathcal{M}$ is a singleton; indeed, (20) is the extension of multiplier preferences to account for multiple models proposed by Hansen and Sargent (2022) and axiomatised by Cerreia-Vioglio et al. (2020). To this extent, confidence aggregation, embedded in the confidence decision framework, can recover several classes of misspecification-motivated decision models. Note that the confidence approach fully separates the epistemic issue of the beliefs (and confidence in them) that can or should be formed on the basis of a set of models from the pragmatic question of their role—as well as that of caution or

[25] Hansen and Sargent's proof relies on the Lagrange multiplier theorem (Luenberger, 1969), and hence on the convexity of $R(q\|m)$ as a function of q. For convex, compact $\mathcal{M}$, $\min_{m\in\mathcal{M}} R(q\|m)$ is a convex function of q, so, for any decision problem in which the standard constraint and multiplier preferences yield the same optima, the same holds for (19) and (20).

ambiguity attitude—in decision making. Misspecification-motivated decision models, by contrast, bake both issues together into the decision rule.

7.3 Performance, expertise and weights

An oft-discussed question related to linear pooling or similar rules concerns the appropriate weights. Influential approaches in risk analysis (Cooke, 1991; Cooke and Goossens, 2008) and psychology (Collins et al., 2023) have developed methods of assigning weights to individuals based on, say, their performance on related questions for which correct replies are available. There is also evidence that such performance-weighted probability aggregation rules outperform equal-weight rules, at least in certain contexts, including in the field of expert judgement (Colson and Cooke, 2018; Budescu and Chen, 2015).

The general message of Section 5 carries over to these approaches: confidence aggregation, by accommodating within-person expertise diversity, will outperform performance-weighted linear pooling in a range of situations. This can be seen on a variant of the case considered in that section.[26] Whilst details differ, a typical performance-weighted approach assigns weights to experts on the basis of past or 'calibration' tasks, under the tacit assumption that they are indicative of performance in the future or on the questions of interest. Using the notation from Section 5, appropriate weights for linear pooling $(v^1, \ldots, v^n)$ can thus be assumed to satisfy: for all individuals $i, i' = 1, \ldots, n$: $v^i \geqslant v^{i'}$ if and only if $\sum_{j=1}^{m} \mathbb{E}\mathbf{p^i}_j \geqslant \sum_{j=1}^{m} \mathbb{E}\mathbf{p^{i'}}_j$. In other words, the higher weights are assigned to those individuals with the better performance, as measured by the closeness to the truth, taken over all the issues. We call issue-independent weights satisfying this property *calibrated*.

If 'calibration' questions can be used to gauge experts' mean performance over all issues, then there may be situations in which they can gauge performance on each issue separately. This can be used to propose issue-specific weights for each individual, for use in the expertise-sensitive pooling rule derived from confidence aggregation (Section 3.3). Analogously to the weights taken for linear pooling above, one can thus set weights w^i_j such that, for all individuals $i, i' = 1, \ldots, n$ and issues $j, j' = 1, \ldots, m$, $w^i_j \geqslant w^{i'}_{j'}$ if and only if $\mathbb{E}\mathbf{p^i}_j \geqslant \mathbb{E}\mathbf{p^{i'}}_{j'}$. In other words, the expert and issue where there is better performance get higher weight. We call issue-specific weights satisfying this property *calibrated*.

Using such weights, one can define performance-weighted linear pooling and expertise-sensitive pooling analogously to the rules considered in Section 5:

[26] In this discussion, all notation is as in Section 5. In particular, we retain the convention that all issues are true.

$$\lambda^{pw}(\mathbf{p^1},\ldots,\mathbf{p^n}) = \sum_{i=1}^{n} v^i \mathbf{p^i} \tag{21}$$

$$\phi^{pw}(\mathbf{p^1},\ldots,\mathbf{p^n}) = \arg\min_{\mathbf{x}} \sum_{i=1}^{n} (\mathbf{x}-\mathbf{p}^i)^T {\mathbf{\Upsilon^i}}^{-1} (\mathbf{x}-\mathbf{p}^i) \tag{22}$$

where $\mathbf{\Upsilon^i}$ is the diagonal matrix with jj^{th} entry w^i_j.

As shown in Appendix B.3 (Proposition 6), under some mild conditions on the relationship between the issue-specific and issue-independent weights, similar results to Theorem 4 hold here: the performance-weighted expertise-sensitive pooling rule ϕ^{pw} outperforms the corresponding linear pooling rule λ^{pw}, in terms of closeness to the truth. Whilst linear pooling is taken as the reference in the cited literatures (Cooke, 1991), some approaches have dabbled with issue-specific weights (Collins et al., 2023). The developments in this paper provide further support for such approaches, beyond proposing a coherent, well-defined method for incorporating issue-specific expertise that overcomes known challenges, as discussed in Section 3.3.

Whilst this paper focuses on aggregation of non-categorical beliefs—i.e. probabilities or generalisations thereof—there is of course a large diverse literature on the aggregation of categorical, yes-or-no judgements, parts of which consider performance (e.g. Condorcet, 1785; Galton, 1907; Surowiecki, 2005). For instance, Prelec et al. (2017) proposes an algorithm for aggregating categorical judgements which relies on participants' judgements about others' responses and provides evidence that, in a variety of settings, it outperforms more traditional methods, including linear pooling (which, confusingly from the perspective of the current paper, they call 'confidence-weighted voting'). Examination of consequences of the confidence-in-belief-based approach developed here in such contexts is left as a topic for future research.

References

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

Aliprantis, C. D. and Border, K. C. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, Berlin, 3rd edition.

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.

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., and Frith, C. D. (2010). Optimally interacting minds. *Science (New York, N.Y.)*, 329(5995):1081–1085.

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

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3.

Budescu, D. V. and Chen, E. (2015). Identifying expertise to extract the wisdom of crowds. *Management Science*, 61(2):267–280. Publisher: INFORMS.

Cerreia-Vioglio, S., Hansen, L. P., Maccheroni, F., and Marinacci, M. (2020). Making decisions under model misspecification. *arXiv preprint arXiv:2008.01071*.

Collins, R. N., Mandel, D. R., and Budescu, D. (2023). Performance-weighted aggregation: ferreting out wisdom within the crowd.

Colson, A. R. and Cooke, R. M. (2018). Expert Elicitation: Using the Classical Model to Validate Experts' Judgments. *Review of Environmental Economics and Policy*, 12(1):113–132. Publisher: The University of Chicago Press.

Condorcet, M. J. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, volume 252. American Mathematical Soc.

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

Cooke, R. M. and Goossens, L. L. (2008). TU Delft expert judgment data base. *Reliability Engineering & System Safety*, 93(5):657–674. Publisher: Elsevier.

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.

Fleming, S. M. and Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8.

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

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

Galton, F. (1907). Vox Populi. *Nature*, 75(1949):450–451.

Genest, C. and Zidek, J. V. (1986). Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, 1(1):114–135.

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

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

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

Green, D. M. and Swets, J. A. (1966). *Signal detection theory and psychophysics*, volume 1. Wiley New York.

Hansen, L. P. (2007). Beliefs, Doubts and Learning: Valuing Macroeconomic Risk. *The American Economic Review*, 97(2):1–30.

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

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

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

Hardy, G., Littlewood, J. E., and Polya, G. (1934). *Inequalities*. Cambridge University Press, Cambridge.

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological review*, 18(3):186. Publisher: The Review Publishing Company.

Hill, B. (2012). Unanimity and the aggregation of multiple prior opinions. GREGHEC Research Papers 959, HEC Paris.

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

Hill, B. (2015). Confidence as a Source of Deferral. Technical Report HEC Paris Research Paper No. ECO/SCD-2014-1060, HEC Paris. Mimeo HEC Paris.

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. ISSN: 2640-3498.

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. Publisher: Elsevier.

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.

Koriat, A. (2012a). The self-consistency model of subjective confidence. *Psychological review*, 119(1):80.

Koriat, A. (2012b). When are two heads better than one and why? *Science*, 336(6079):360–362. Publisher: American Association for the Advancement of Science.

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

Mamassian, P. (2016). Visual confidence. *Annual Review of Vision Science*, 2(1):459–481. Publisher: Annual Reviews.

Marshall, A. W., Olkin, I., and Arnold, B. C. (2011). *Inequalities: Theory of Majorization and Its Applications*. Springer Series in Statistics. Springer, New York, NY.

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

Metcalfe, J. and Shimamura, A. P. (1994). *Metacognition: Knowing about knowing*. MIT press.

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.

Nascimento, L. (2012). The Ex-Ante Aggregation of Opinions under Uncertainty. *Theoretical Economics*.

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

Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychological review*, 117(3):864. Publisher: American Psychological Association.

Prelec, D., Seung, H. S., and McCoy, J. (2017). A solution to the single-question crowd wisdom problem. *Nature*, 541(7638):532–535. Publisher: Nature Publishing Group UK London.

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.

Rahnev, D., Desender, K., Lee, A. L., Adler, W. T., Aguilar-Lleyda, D., Akdoğan, B., Arbuzova, P., Atlas, L. Y., Balcı, F., and Bang, J. W. (2020). The confidence database. *Nature human behaviour*, 4(3):317–325. Publisher: Nature Publishing Group.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological review*, 85(2):59. Publisher: American Psychological Association.

Ratcliff, R. and Starns, J. J. (2013). Modeling confidence judgments, response times, and multiple choices in decision making: recognition memory and motion discrimination. *Psychological review*, 120(3):697. Publisher: American Psychological Association.

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.

Surowiecki, J. (2005). *The wisdom of crowds*. Anchor.

Winkler, R. L. and Murphy, A. H. (1968). “Good” probability assessors. *Journal of Applied Meteorology and Climatology*, 7(5):751–758.

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

Table 2: Characterisations of special cases

A Characterising confidence aggregation: special cases

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

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 $\bigcup_{\mathbf{o}:\otimes\mathbf{o}=o} \bigcap_i c'^i(o_i) \neq \varnothing$

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

The only difference with respect to the representation involved in Theorem 3 is that here the union is taken over all tuples of confidence levels whose confidence-level aggregate equals o, whereas the previous procedure looks at all those with confidence-level aggregate at most o. It follows directly from the nestedness property of confidence rankings (i.e. the fact that c is 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 1.

We have the following result, which involves the axioms in Figure 4, and defines clauses according to Table 2.

Theorem 6. *Suppose that O is infinite, and let $\{\succ^i\}, \succ^0$ satisfy Assumptions 1 and 2. For each of the rows in Table 2: $\{\succ^i\}, \succ^0$ 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 2 on normative grounds; we present this result to illustrate the richness of the approach, and exemplify some simple aggregators.

The axiom involved in the characterisation of confidence aggregation with an affine aggregator, Consensus Independence, uses the notion of uncovered consensus. For every tuple of stakes levels $\mathbf{s}$ exhibiting consensus 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 $\mathbf{s'} \not\geqslant \mathbf{s}$ 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 the required preference. 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, and 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 that the mixture of uncovered consensuses can 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

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 (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 $\alpha \in [0,1]$, $(\alpha\mathbf{s} + (1-\alpha)\mathbf{t})_i = \alpha s_i + (1-\alpha) t_i$ and $(\mathbf{s} \vee \mathbf{t})_i = \max\{s_i, t_i\}$.

Figure 4: Axioms for special cases

consensus corresponding to the largest stakes level for each entry in the tuple (the join). Neutrality is a standard neutrality axiom, adapted to the current context, stating that respect for consensus is preserved under permutation of individuals. Added to each of the other conditions, it characterises the 'neutral' version of affine and generalised maximum aggregators respectively.

B Proofs

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

We begin with the following Proposition, mentioned in Section 2.

Proposition 3. *c^0 is a consensus-preserving aggregation of $(c^1, \ldots, c^n)$ under $\otimes$ if and only if*

$$\iota^0(p) = \otimes(\iota^1(p), \ldots, \iota^n(p))$$

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. $p \in c^i(\iota^i(p))$ 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 required formula follows from the definition of ι. □

We now prove Theorems 3 and 6. 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.

B.1.1 Proof of Theorem 3

We first show sufficiency of the axioms. Recall that $\{(c^i, D, u)\}, (c^0, D, u)$ denote the representations of the $\{\succ^i\}, \succ^0$. 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 consistent; but $\succ_{t_i} \subseteq \succ_{s_i}$ for all i by the properties of confidence rankings, so $\succ_\mathbf{t}$ is consistent and hence $\mathbf{t} \in X$.

We say that a preference relation $\succ$ is a *Bewley preference* if there exists a representation à la Bewley: i.e. there exists a closed convex set of priors $\mathcal{P} \subseteq \Delta$ and utility function u' such that, for every $f, g \in \mathcal{A}$, $f \succ g$ if and only if

$$\mathbb{E}_p u'(f) > \mathbb{E}_p u'(g) \quad \text{for all } p \in \mathcal{P} \tag{B.1}$$

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

Claim 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 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, $f \succ_{\mathbf{s}} g$ if and only if

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

Proof. First consider $\succ_{\mathbf{s}}$ exhibiting consensus, and let $\succ_{\bigcap \mathbf{s}}$ be the Bewley preference with utility u and set of priors $\bigcap_{i=1}^n c^i(s_i)$. 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 $f \nprec^i_{s_i} g$ 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:

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

where the equality follows from Claim 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 4. *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 3. For every $s \geqslant \min G(X)$, $\succ^0_s$ 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 $f, g \in \mathcal{A}$, $f \succ^0_s g$ if and only if:

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

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 $\succ^0_{s''} \subseteq \succ_{\mathbf{s}'}$. It follows from the nestedness properties of confidence rankings that $\succ^0_s \subseteq \succ^0_{s''} \subseteq \succ_{\mathbf{s}'}$. Since this holds for all $\mathbf{s}'$ with $G(\mathbf{s}') \leqslant s$, it follows that $\succ^0_s \subseteq \bigcap_{\mathbf{s}\in X: s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$.

To establish the opposite containment, consider f, g with $f \succ_{\mathbf{s}} g$ for all $\mathbf{s} \in X$ with $s \geqslant G(\mathbf{s})$. For any $\mathbf{s}'$ such that $\succ^0$ respects the consensus $\succ_{\mathbf{s}'}$ at s, it follows from the definition of G that $s \geqslant G(\mathbf{s}')$, so $f \succ_{\mathbf{s}'} g$ by the assumption specifying f, g. Hence, by Corpus-wise Pareto, $f \succ^0_s g$. So $\succ^0_s \supseteq \bigcap_{\mathbf{s}\in X: s \geqslant G(\mathbf{s})} \succ_{\mathbf{s}}$, and hence there is equality. It follows from Claim 2 that (B.3) holds for all $s \geqslant \min G(X)$.

□

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

By the nestedness of confidence rankings, 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

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

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 ($G(\mathbf{s}) = G(\mathbf{t})$). Under the simplifying assumption that ζ is the identity, $\otimes = G$ on X. Moreover, $\otimes$ is monotonic, and thus a confidence level aggregator. It follows from Claim 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 1 and 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 2. 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 (A.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 $\otimes\mathbf{o}' = o$. The following is an example where this condition is not satisfied.

Example B.1. 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.

B.1.2 Proof of Theorem 6

Proof of part i. (affine aggregation). Let X be as defined in the proof of Theorem 3. Let

$$C = \{(\mathbf{s}, s) \in \mathbb{R}^{n+1} : \mathbf{s} \in X, \ \succ^0_s \subseteq \succ_{\mathbf{s}}\}$$
$$K = \{(\mathbf{s}, s) \in C : \succ_{\mathbf{s}} \supseteq \bigcap_{\mathbf{s}' : (\mathbf{s}', s) \in C, \ \mathbf{s}' \ngeqslant \mathbf{s}} \succ_{\mathbf{s}'}\}$$

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

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

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 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 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 3. By Proposition 4 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 6. For all $s \geqslant \inf G_{aff}(X)$, $\succ^0_s$ is represented by $\bigcup_{\mathbf{s}\in X:s=G_{aff}(\mathbf{s})} \bigcap_i c^i(s_i)$ in the Bewley sense: i.e. for all $f, g \in \mathcal{A}$, $f \succ^0_s g$ if and only if:

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

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, $\succ^0_s \subseteq \succ_{\mathbf{s}}$. So $\succ^0_s \subseteq \bigcap_{\mathbf{s}\in X:s=G_{aff}(\mathbf{s})} \succ_{\mathbf{s}}$.

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

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

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

Define $\otimes$ by

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

Clearly, this is an affine confidence level aggregator. Moreover, by Claim 6 and the fact that $\zeta : \mathcal{S} \to O$ is the identity, (A.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 3, 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 $c^0\left(\sum_i w_i(\mathbf{s_k})_i + \chi\right) \supseteq \bigcap_i c^i((\mathbf{s_k})_i)$, 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 $\bigcap_i c^i((\mathbf{t_k})_i) \nsubseteq \bigcup_{\mathbf{s} \not\geqslant \mathbf{t_k},\ (\mathbf{s},t_k) \in C} \bigcap_i c^i(s_i) \subseteq c^0(t_k)$. By the aggregation representation, it follows that $c^0(\sum_i w_i (\mathbf{t_k})_i + \chi) = \bigcup_{\mathbf{s}:\ \sum_i w_i s_i = \sum_i w_i (\mathbf{t_k})_i} \bigcap_i c^i(s_i) \subseteq \bigcap_i c^i((\mathbf{t_k})_i) \cup \bigcup_{\mathbf{s} \not\geqslant \mathbf{t_k},\ (\mathbf{s},t_k) \in C} \bigcap_i c^i(s_i) \subseteq c^0(t_k)$, so, by the nestedness of confidence rankings, $\sum_i w_i (\mathbf{t_k})_i + \chi \leqslant t_k$. So if $\succ_{\mathbf{t_k}}$ is an uncovered consensus at t_k for each k, it follows that $\sum_{k=1}^{m} \beta_k t_k \geqslant \sum_i w_i \sum_{k=1}^{m} \beta_k (\mathbf{t_k})_i + \chi$. Since, $\sum_i w_i \sum_{k=1}^{l} \alpha_k (\mathbf{s_k})_i = \sum_i w_i \sum_{k=1}^{m} \beta_k (\mathbf{t_k})_i$, it follows that $\sum_{k=1}^{m} \beta_k t_k > \sum_{k=1}^{l} \alpha_k s_k$, as required.

□

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

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

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

Now consider the case where there does not exist s and $\mathbf{s}$ such that $(\mathbf{s}, s) \in C$, $\mathbf{s}$ is a maximum, under $\geqslant$, of $\{\mathbf{s}' : (\mathbf{s}', s) \in C\}$, and $s_j \neq s_k$. Hence, for all s and $\mathbf{s}$ such that $(\mathbf{s}, s) \in C$ and $\mathbf{s}$ is a maximum, under $\geqslant$, of $\{\mathbf{s}' : (\mathbf{s}', s) \in C\}$, $s_j = s_k$. For any $\mathbf{s}$, let $\hat{\mathbf{s}}$ be such that: $\hat{s}_i = s_i$ when $i \neq j, k$, $\hat{s}_j = \hat{s}_k = \max\{s_j, s_k\}$. So, in the case under consideration, for every $\mathbf{s}$ with $s_j \neq s_k$ and every stakes level s, $(\mathbf{s}, s) \in C$ if and only if

$(\hat{\mathbf{s}}, s) \in C$.

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

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

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

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

up to convex closure.

For any $\mathbf{s}$ 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 of part iii. (generalised maximum aggregator). Consider G as defined in the proof of Theorem 3. 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 $s_i > t_i^s$ for some i, $G(\mathbf{s}) > s$. Since, for any $\mathbf{s}$, if $\mathbf{s} \leqslant \mathbf{t}^s$, then $G(\mathbf{s}) \leqslant s$ by the monotonicity of G, we have that, for all $\mathbf{s}$, $G(\mathbf{s}) > s$ if and only if there exists i with $s_i > t_i^s$. Hence $G(\mathbf{s}) < s$ if and only if there exists $s' < s$ with $s_i \leqslant t_i^{s'}$ for all i. Hence $G(\mathbf{s}) = s$ if and only if $\mathbf{s} \leqslant \mathbf{t}^s$ and there is no $s' < s$ with $\mathbf{s} \leqslant \mathbf{t}^{\mathbf{s}'}$.

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

For $i = 1, \ldots, n$, define $\psi_i : O \to O$ by $\psi_i(o) = \min\{s : t_i^s \geqslant o\}$. This is well-defined, since $\zeta : \mathcal{S} \to O$ is the identity map. Since, by the confidence representation, t_i^s 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 $\mathbf{s} \leqslant \mathbf{t^s}$ and $\mathbf{s} \nleqslant \mathbf{t^{s'}}$ 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 would be even if ζ were not invertible), 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 $\bigcup_{\mathbf{o}:\otimes o_i = o} \bigcap_i c^i(o_i) \neq \varnothing$, (A.1) holds with $\otimes$, up to convex closure. By a similar argument to that used in the proof of Theorem 3, 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 of part iv. (maximum aggregator). Consider $\mathbf{t^s}$, as defined in the proof of part iii; we show that $t^s_j = t^s_k$ for all j, k. For reductio, suppose that this is not the case for some j, k, and suppose without loss of generality that $t^s_j > t^s_k$. By Neutrality, $G((t^s_k)_j(t^s_j)_k\mathbf{t^s}) = G(\mathbf{t^s}) = s$; but since $t^s_j > t^s_k$, it follows by the properties of G established in the proof of part iii. that $G((t^s_k)_j(t^s_j)_k\mathbf{t^s}) > G(\mathbf{t^s}) = s$, which is a contradiction. So $t^s_j = t^s_k$ 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.

□

B.2 Proofs of results in Section 3

Proof of Theorem 1. By (3), the centre of c is:

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

In part i., $\iota^i(p) = w^i \sum_{\omega\in\Omega'}(p(\omega) - p_i(\omega))^2$, 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.

In part ii., $\iota^i(p) = w^i R(p\|p_i)$, 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.

In part iii., $\iota^i(p) = w^i R(p_i\|p)$, 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 5. *Under the conditions and setup of Example 3.1, let c^{sq} be the w^L Euclidean confidence ranking generated by p^L (with $\omega' = \omega_R$), c^{rKL} be the w^L reverse relative entropy confidence ranking generated by p^L. For $\epsilon \in [0, 0.9]$ let $\mathcal{L} = \{p \in \Delta : p(L) \geqslant 0.9 - \epsilon\}$, $\mathcal{R} = \{p \in \Delta : p(L) \leqslant 0.1 + \epsilon\}$. Then, for all $o \in O$, $c^{sq}(o) \subseteq \mathcal{L}$ if and only if $c^{sq}(o) \subseteq \mathcal{R}$, and $c^{rKL}(o) \subseteq \mathcal{L}$ if and only if $c^{rKL}(o) \subseteq \mathcal{R}$.*

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 $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), p($
More specifically, in the Euclidean case, with $\omega' = \omega_R$,

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

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 $p(L) = p(\omega_L) + p(\omega_{LR}) = 0.6$. 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 $\phi(p^L(\omega_{LR}), p^L(\omega_L), p^L(\omega_N), q'(\omega_{LR}), q'(\omega$ 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 1. Take $o = w^L_L \epsilon^2 + w^L_B (\max\{\epsilon - 0.81, 0\})^2$. q, defined by $q(\omega_{LR}) = p^L(\omega_{LR}) - \max\{\epsilon - 0.81, 0\} = 0.09 - \max\{\epsilon - 0.81, 0\}$, $q(\omega_R) = p^L(\omega_R) + \max\{\epsilon - 0.81, 0\} = 0.01 + \max\{\epsilon - 0.81, 0\}$, $q(\omega_L) = p^L(\omega_L) - \epsilon = 0.81 - \min\{\epsilon, 0.81\}$ 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, $\sum_{j=\{L,R,B\}} w^L_j (q(j) - p^L(j))^2 = w^L_L \epsilon^2 + w^L_B (\max\{\epsilon - 0.81, 0\})^2$, so $q \in c^L(o)$. Since, for any q' with $q'(L) < 0.9 - \epsilon$, $\sum_{j=\{L,R,B\}} w^L_j (q'(j) - p^L(j))^2 > w^L_L \epsilon^2 + w^L_B (\max\{\epsilon - 0.81, 0\})^2$, such $q' \notin c^L(o)$, so $c^L(o) \subseteq \mathcal{L}$. 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. $\sum_{j=\{L,R,B\}} w^L_j (q_\delta(j) - p^L(j))^2 = w^L_R \delta^2 + w^L_B (\max\{0, \delta - 0.09\})^2$. Noting that $w^L_R \epsilon^2 + w^L_B (\max\{0, \epsilon - 0.09\})^2 < w^L_L \epsilon^2 + w^L_B (\max\{\epsilon - 0.81, 0\})^2$ if and only if $w^L_B \frac{1}{\epsilon^2} \left((\max\{0, \epsilon - 0.09\})^2 - (\max\{\epsilon - 0.81, 0\})^2 \right) < w^L_L - w^L_R$, it is straightforward to check that this is the case for all $\epsilon \in [0, 0.9]$ whenever $0.8 w^L_B = w^L_B \frac{0.81^2 - 0.09^2}{0.9^2} < w^L_L - w^L_R$. It follows that there exists $\delta > \epsilon$ with $\sum_{j=\{L,R,B\}} w^L_j (q_\delta(j) - p^L(j))^2 \leqslant w^L_L \epsilon^2 + w^L_B (\max\{\epsilon - 0.81, 0\})^2 = o$, so $c^L(o) \nsubseteq \mathcal{R}$, as required.

□

B.3 Proofs of results in Section 5

To prove Theorem 4, we prove a stronger theorem, involving the following assumptions.

First, we split Assumption 3 into two assumptions. On the one hand, we retain the assumption that the stochastic processes underlying different individuals' reports are independent.

Assumption 5. *For all $i, i' \in \{1, \ldots, n\}$ with $i \neq i'$, $\mathbf{p^i}$ and $\mathbf{p^{i'}}$ are independent (so $\mathbb{E}(\mathbf{p^i}\mathbf{p^{i'}}) = \mathbb{E}(\mathbf{p^i})\mathbb{E}(\mathbf{p^{i'}})$).*

On the other hand, we no longer assume that all issues are independent, i.e. that the covariance matrices $\mathbf{\Gamma^i}$ and $\mathbf{\Sigma^i}$ are diagonal. Rather, for each i, since $\mathbf{\Gamma^i}$ is a positive-definite covariance matrix, there exist an orthonormal matrix $\mathbf{P^i}$, constructed from the principal components of $\mathbf{\Gamma^i}$, and a diagonal matrix $\mathbf{E^i}$ with $\mathbf{\Gamma^i} = \mathbf{P^i}\mathbf{E^i}\mathbf{P^i}^T$. Note that $\mathbf{P^i}$ can always be chosen so that $(\mathbf{P^i}^T\mathbf{1})_j \geqslant 0$ for each j; henceforth we assume that this holds for all $\mathbf{P^i}$. Since $\mathbf{\Sigma^i}$ is also positive-definite, there is a similar decomposition for it: $\mathbf{\Sigma^i} = \mathbf{Q^i}\mathbf{D^i}\mathbf{Q^i}^T$ for a diagonal matrix $\mathbf{D^i}$. We replace independence across issues with the following simplifying assumption throughout:

Assumption 6. *For all i, i', $\mathbf{Q^i} = \mathbf{P^i} = \mathbf{P^{i'}} = \mathbf{Q^{i'}}$.*

This says that, whatever 'distortion' there is between the true probability distribution over answers and what is reflected in the confidence, it does not change the principal components reflecting a subjects' minimal view of the correlation across issues. Moreover, the individuals share a minimum view of the correlation across issues, in the sense that their covariance matrices share the same principal components. Recall that the variability in probability estimates concerning a principal component (a column in $\mathbf{P}$) is entirely captured by the variance of the underlying distribution as concerns that issue—$\mathbf{E^i}_{jj}$. Similarly, the confidence concerning the component is entirely captured by the corresponding element of the diagonal matrix $\mathbf{D^i}_{jj}$. Note that, in the special case in which $\mathbf{\Sigma^i}$ and $\mathbf{\Gamma^i}$ are diagonal (the draws for different events are independent), $\mathbf{D^i}_{jj}$ and $\mathbf{E^i}_{jj}$ are just the appropriate variances. Clearly, Assumption 3 implies Assumptions 5 and 6 (with $\mathbf{P}$ the identity matrix in the latter case).

In the light of this assumption, Assumption 4 needs to be modified to apply on the common principal components, as follows:

Assumption 7. *For all i, i', j, $\mathbf{E^i}_{jj} \geqslant \mathbf{E^{i'}}_{jj}$ if and only if $\mathbf{D^i}_{jj} \geqslant \mathbf{D^{i'}}_{jj}$ if and only if $(\mathbf{P}^T(\mathbf{1}-\mu^{\mathbf{i}}))_j \geqslant (\mathbf{P}^T(\mathbf{1}-\mu^{\mathbf{i'}}))_j$.*

Again, when $\mathbf{P}$ is the identity, as implied by Assumption 3, this coincides with Assumption 4.

We have the following Theorem, of which Theorem 4 is clearly a direct corollary.

Theorem 7. *Under Assumptions 5, 6 and 7:*

i in expectation, the ℓ_1 distance from the truth is smaller under confidence aggregation with the average confidence-level aggregator, as compared to linear pooling:

$$\mathbb{E}\sum_{j=1}^{m}\left(\mathbf{1}-\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))\right)_j \leqslant \mathbb{E}\sum_{j=1}^{m}\left(\mathbf{1}-\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))\right)_j$$

ii in expectation, the Brier score is smaller under confidence aggregation with the average confidence-level aggregator, as compared to linear pooling:

$$\mathbb{E}B(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))) \leqslant \mathbb{E}B(\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))$$

Proof of Theorem 7 . Let $\mathbf{A^i} = (\mathbf{\Sigma^i})^{-1}$, $\mathbf{E^i}_{jj} = (\sigma^i_j)^2$ and $\mathbf{D^i}_{jj} = (\rho^i_j)^2$.Then the centre of the confidence ranking is $\mathbf{x}$ minimising:

$$\sum_{i=1}^{n}(\mathbf{x}-\mathbf{p}^i)^T\mathbf{A^i}(\mathbf{x}-\mathbf{p}^i)$$

Differentiating, we get the FOC for the solution $\mathbf{x}^*$:

$$2\sum_{i=1}^{n}\mathbf{A^i}(\mathbf{x}^*-\mathbf{p}^i) \quad = \quad 0 \tag{B.5}$$

hence

$$\left(\sum_{i=1}^{n}\mathbf{A^i}\right)\mathbf{x}^* = \sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i \tag{B.6}$$

Taking the expectations, we have

$$\left(\sum_{i=1}^{n}\mathbf{A^i}\right)\mathbb{E}\mathbf{x}^* = \sum_{i=1}^{n}\mathbf{A^i}\mu^{\mathbf{i}} \tag{B.7}$$

Moreover, noting that the previous equalities imply that:

$$\left(\sum_{i=1}^{n}\mathbf{A^i}\right)\mathbf{x}^* - \mathbb{E}\left[\left(\sum_{i=1}^{n}\mathbf{A^i}\right)\mathbf{x}^*\right] = \sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i - \mathbb{E}\left[\sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i\right] \tag{B.8}$$

we have, by multiplying with their own transpose and taking the expectation, that:

$$\begin{aligned} &\left(\sum_{i=1}^{n}\mathbf{A^i}\right)\mathbb{E}\left[(\mathbf{x}^*-\mathbb{E}\mathbf{x}^*)(\mathbf{x}^*-\mathbb{E}\mathbf{x}^*)^t\right]\left(\sum_{i=1}^{n}\mathbf{A^i}\right)^T \\ = \;& \mathbb{E}\left(\left(\sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i - \mathbb{E}\left[\sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i\right]\right)\left(\sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i - \mathbb{E}\left[\sum_{i=1}^{n}\mathbf{A^i}\mathbf{p}^i\right]\right)^T\right) \end{aligned} \tag{B.9}$$

By Assumption 6, $\mathbf{A^i} = \mathbf{P}\mathbf{D^i}^{-1}\mathbf{P}^T$, so $\sum_{i=1}^{n}\mathbf{A^i} = \mathbf{P}\left(\sum_{i=1}^{n}\mathbf{D^i}^{-1}\right)\mathbf{P}^T$; henceforth let $\hat{\mathbf{D}} = \sum_{i=1}^{n}\mathbf{D^i}^{-1}$. From (B.7), we have that

$$\begin{aligned} \mathbf{P}\hat{\mathbf{D}}\mathbf{P}^T\mathbb{E}\mathbf{x}^* &= \sum_{i=1}^{n}\mathbf{P}\mathbf{D^i}^{-1}\mathbf{P}^T\mu^{\mathbf{i}} \\ \mathbf{P}^T\mathbb{E}\mathbf{x}^* &= \sum_{i=1}^{n}\hat{\mathbf{D}}^{-1}\mathbf{D^i}^{-1}\mathbf{P}^T\mu^{\mathbf{i}} \end{aligned}$$

It follows that:

$$(\mathbf{P}^T\mathbb{E}\mathbf{x}^*)_j = \frac{1}{\sum_{i=1}^{n}\frac{1}{(\rho_j^i)^2}}\sum_{i=1}^{n}\frac{1}{(\rho_j^i)^2}(\mathbf{P}^T\mu^{\mathbf{i}})_j$$

for each $j = 1,\ldots,m$. By Assumption 7, $\frac{1}{(\rho_j^i)^2} \leqslant \frac{1}{(\rho_j^{i'})^2}$ whenever $(\mathbf{P}^T(\mathbf{1}-\mu^{\mathbf{i}}))_j \geqslant (\mathbf{P}^T(\mathbf{1}-\mu^{\mathbf{i'}}))_j$, so, for all j,

$$(\mathbf{P}^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*))_j \leqslant (\mathbf{P}^T(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}}))_j \tag{B.10}$$

Hence $\left(\mathbf{P}^T\mathbb{E}\mathbf{x}^*\right)_j - \left(\sum_{i=1}^{n}\frac{1}{n}\mathbf{P}^T\mu^{\mathbf{i}}\right)_j \geqslant 0$ for all j. Since $(\mathbf{P}^T\mathbf{1})_j \geqslant 0$ for all j, it follows that

$$\mathbb{E}\sum_{j=1}^{m}\big(\mathbf{1}-\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))\big)_j-\mathbb{E}\sum_{j=1}^{m}\big(\mathbf{1}-\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))\big)_j \tag{B.11}$$

$$\begin{aligned}
&=\sum_{j=1}^{m}(\mathbf{1}-\mathbb{E}\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))_j-\sum_{j=1}^{m}(\mathbf{1}-\mathbb{E}\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,\mathbf{p^n},\mathbf{\Sigma^n})))_j\\
&=\mathbf{1}^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*)-\mathbf{1}^T\left(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}}\right)\\
&=\mathbf{1}^T\mathbf{P}\left(\mathbf{P}^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*)\right)-\mathbf{1}^T\mathbf{P}\left(\mathbf{P}^T\left(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}}\right)\right)\\
&\leqslant 0
\end{aligned} \tag{B.12}$$

establishing clause i. of the Theorem.

As concerns clause ii., first note that, from (B.10), we have:

$$\begin{aligned}
&\sum_{j=1}^{m}(1-\mathbb{E}\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2-\sum_{j=1}^{m}(1-\mathbb{E}\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))_j)^2\\
&=(\mathbf{1}-\mathbb{E}\mathbf{x}^*)^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*)-\left(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}}\right)^T\left(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}}\right)\\
&=(\mathbf{P}^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*))^T\left(\mathbf{P}^T(\mathbf{1}-\mathbb{E}\mathbf{x}^*)\right)-\left(\mathbf{P}^T(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}})\right)^T\left(\mathbf{P}^T(\mathbf{1}-\sum_{i=1}^{n}\frac{1}{n}\mu^{\mathbf{i}})\right)\\
&\leqslant 0
\end{aligned} \tag{B.13}$$

Moreover, by (B.9), we have the following expression for the covariance matrix of $\mathbf{x}^*$

$$\begin{aligned}
\mathbf{K}_{\mathbf{x}^*\mathbf{x}^*} &= \mathbb{E}\left[(\mathbf{x}^* - \mathbb{E}\mathbf{x}^*)(\mathbf{x}^* - \mathbb{E}\mathbf{x}^*)^T\right] \\
&= \left(\sum_{i=1}^n \mathbf{A^i}\right)^{-1} \left[\mathbb{E}\left(\left(\sum_{i=1}^n \mathbf{A^i}\mathbf{p}^i - \mathbb{E}\left[\sum_{i=1}^n \mathbf{A^i}\mathbf{p}^i\right]\right)\left(\sum_{i=1}^n \mathbf{A^i}\mathbf{p}^i - \mathbb{E}\left[\sum_{i=1}^n \mathbf{A^i}\mathbf{p}^i\right]\right)^T\right)\right]\left(\left(\sum_{i=1}^n \mathbf{A^i}\right)^T\right)^{-1} \\
&= \left(\sum_{i=1}^n \mathbf{A^i}\right)^{-1} \mathbb{E}\left(\left(\sum_{i=1}^n \mathbf{A^i}\left(\mathbf{p}^i - \mathbb{E}\mathbf{p}^i\right)\right)\left(\sum_{i=1}^n \mathbf{A^i}\left(\mathbf{p}^i - \mathbb{E}\mathbf{p}^i\right)\right)^T\right)\left(\left(\sum_{i=1}^n \mathbf{A^i}\right)^T\right)^{-1} \\
&= \left(\sum_{i=1}^n \mathbf{A^i}\right)^{-1} \mathbb{E}\left(\sum_{i=1}^n\sum_{k=1}^n \mathbf{A^i}\left(\mathbf{p}^i - \mathbb{E}\mathbf{p}^i\right)\left(\mathbf{p}^k - \mathbb{E}\mathbf{p}^k\right)^T \mathbf{A^k}^T\right)\left(\left(\sum_{i=1}^n \mathbf{A^i}\right)^T\right)^{-1} \\
&= \left(\sum_{i=1}^n \mathbf{A^i}\right)^{-1} \sum_{i=1}^n \mathbf{A^i}\mathbf{\Gamma^i}\mathbf{A^k}^T\left(\left(\sum_{i=1}^n \mathbf{A^i}\right)^T\right)^{-1} \\
&= \left(\sum_{i=1}^n \mathbf{A^i}\right)^{-1} \sum_{i=1}^n \mathbf{P}\mathbf{D^i}^{-1}\mathbf{E^i}\mathbf{D^i}^{-1}\mathbf{P}^T\left(\left(\sum_{i=1}^n \mathbf{A^i}\right)^T\right)^{-1} \\
&= \mathbf{P}\hat{\mathbf{D}}^{-1}\left(\sum_{i=1}^n \mathbf{D^i}^{-1}\mathbf{E^i}\mathbf{D^i}^{-1}\right)\hat{\mathbf{D}}^{-1}\mathbf{P}^T
\end{aligned}$$

where the fourth equality holds due to Assumption 5, and the last two by Assumption 6. So:

$$\begin{aligned}
\left(\mathbf{P}^T\mathbf{K}_{\mathbf{x}^*\mathbf{x}^*}\mathbf{P}\right)_{jj} &= \frac{1}{\left(\sum_{i=1}^n \frac{1}{(\rho_j^i)^2}\right)^2}\sum_{i=1}^n \frac{(\sigma_j^i)^2}{(\rho_j^i)^4} \\
&= \sum_{i=1}^n (\sigma_j^i)^2\left(\frac{\frac{1}{(\rho_j^i)^2}}{\left(\sum_{i=1}^n \frac{1}{(\rho_j^i)^2}\right)}\right)^2 \\
&\leqslant \sum_{i=1}^n (\sigma_j^i)^2 \sum_{i=1}^n\left(\frac{\frac{1}{(\rho_j^i)^2}}{\left(\sum_{i=1}^n \frac{1}{(\rho_j^i)^2}\right)}\right)^2 \\
&\leqslant \frac{1}{n^2}\sum_{i=1}^n (\sigma_j^i)^2 \qquad \text{(B.14)}
\end{aligned}$$

for all j, where the first inequality holds by the rearrangement inequality and Assumption 7 (which implies that, for all j and individuals i, i', $\frac{1}{(\rho_j^i)^2} \geqslant \frac{1}{(\rho_j^{i'})^2}$ if and only if $(\sigma_j^i)^2 \leqslant (\sigma_j^{i'})^2$), and the second inequality holds by basic mathematics.[27] Hence, using $\sigma^2_{\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j}$ to denote the variance of ${\mathbf{x}^*}_j$[28] and similarly for $\sigma^2_{\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j}$

[27] The maximal value of $\sum_{1=1}^n a_i^2$ for $a_i \geqslant 0$ such that $\sum_{1=1}^n a_i = 1$ is obtained when $a_i = \frac{1}{n}$ for all i.

[28] I.e. $\sigma^2_{\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j} = \mathbb{E}(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j - \mathbb{E}\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2$.

and $\frac{1}{n}\sum_{i=1}^{n}\mathbf{p^i}_j$, we have:

$$\begin{aligned}\sum_{j=1}^{m}\sigma^2_{\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j} =& tr(\mathbf{K_{x^*x^*}})\\ =& tr\left(\mathbf{P}^T\mathbf{K_{x^*x^*}}\mathbf{P}\right)\\ \leqslant& \sum_{j=1}^{m}\frac{1}{n^2}\sum_{i=1}^{n}(\sigma^i_j)^2\\ =& tr\left(\mathbf{P}^T\left(\frac{1}{n^2}\sum_{i=1}^{n}\mathbf{\Gamma^i}\right)\mathbf{P}\right)\\ =& tr\left(\frac{1}{n^2}\sum_{i=1}^{n}\mathbf{\Gamma^i}\right)\\ =& \sum_{j=1}^{m}\sigma^2_{\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j}\end{aligned} \tag{B.15}$$

It follows from (B.12) and (B.15) that, for all j,

$$\begin{aligned}&\mathbb{E}B(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))\\ =&\sum_{j=1}^{m}\mathbb{E}(1-\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2\\ =&\sum_{j=1}^{m}(1-\mathbb{E}(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2\\ &+\sum_{j=1}^{m}\mathbb{E}(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j-\mathbb{E}\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2\\ =&\sum_{j=1}^{m}(1-\mathbb{E}(\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2+\sum_{j=1}^{m}\sigma^2_{\phi((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j}\\ \leqslant&\sum_{j=1}^{m}(1-\mathbb{E}(\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2+\sum_{j=1}^{m}\sigma^2_{\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))j}\\ =&\sum_{j=1}^{m}\mathbb{E}(1-\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n}))_j)^2\\ =&\mathbb{E}B(\lambda((\mathbf{p^1},\mathbf{\Sigma^1}),\ldots,(\mathbf{p^n},\mathbf{\Sigma^n})))\end{aligned}$$

establishing clause ii. of the Theorem.

□

The following Proposition, which can be considered a partial extension of the previous result to different weights, uses the terminology introduced in Section 7.3.

Proposition 6. *Let λ^{pw} and ϕ^{pw} be as defined in equations (21) and (22), with calibrated weights w^i_j and v^i, where $v^i = \frac{1}{\sum_{i=1}^{n}\sum_{j=1}^{m}w^i_j}\sum_{j=1}^{m}w^i_j$ for all $i = 1,\ldots,n$ and $\sum_{i=1}^{n}w^i_j = 1$*

for all $j = 1, \ldots, m$. *Then:*

$$\mathbb{E}\sum_{j=1}^{m}\left(\mathbf{1} - \phi^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})\right)_j \leqslant \mathbb{E}\sum_{j=1}^{m}\left(\mathbf{1} - \lambda^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})\right)_j$$

Moreover, under Assumptions 3 and 4:

$$\mathbb{E}B(\phi^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})) \leqslant \mathbb{E}B(\lambda^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n}))$$

Proof. Let μ be the $\mathbb{R}^{n\times m}$ vector of $\mathbb{E}\mathbf{p^i}_j$, written in increasing order: i.e. there is a bijection $\tau : \{1, \ldots, n \times m\} \to \{1, \ldots, n\} \times \{1, \ldots, m\}$ with $\mu_k = \mathbb{E}\mathbf{p}^{\tau_1(\mathbf{k})}{}_{\tau_2(k)}$ for all $k = 1, \ldots, n \times m$, and, for all k, k', $k > k'$ implies $\mu_k \geqslant \mu_{k'}$. Let σ be the $\mathbb{R}^{n\times m}$ vector of variances, under the same ordering: i.e. $\sigma_k = (\sigma^2)^{\tau_1(k)}_{\tau_2(k)} = \mathbb{E}\left(\mathbf{p}^{\tau_1(\mathbf{k})}{}_{\tau_2(k)} - \mathbb{E}\mathbf{p}^{\tau_1(\mathbf{k})}{}_{\tau_2(k)}\right)^2$ for all $k = 1, \ldots, n\times m$. Under Assumption 4, this vector is in decreasing order. Let $\mathbf{w}$ be the corresponding vector of issue-dependent weights, ordered according to the order of μ: $\mathbf{w}_k = w^{\tau_1(k)}_{\tau_2(k)}$ for all $k = 1, \ldots, n\times m$. Since (w^i_j) are calibrated, this vector is in increasing order: for all k, k', if $k \geqslant k'$, then $\mathbf{w}_k \geqslant \mathbf{w}_{k'}$. Finally, let $\mathbf{v}$ be the corresponding vector of issue-independent weights, $\mathbf{v}_k = v^{\tau_1(k)}$. Note that, since $\sum_i w^i_j = 1$ for all j, it follows that $v^i = \frac{\sum_j w^i_j}{\sum_i \sum_j w^i_j} = \frac{\sum_j w^i_j}{m}$ for each $i = 1, \ldots, n$.

Let $\sqsubset$ denote majorization of series, in the sense of Hardy et al. (1934); Marshall et al. (2011), and let $\sqsubset_w$ denote weak majorization.[29] By standard arguments, for each i:

$$(v^i, \ldots, v^i) = \sum_j w^i_j \left(\frac{1}{m}, \ldots, \frac{1}{m}\right) \sqsubset (w^i_1, \ldots, w^i_m)$$

Whence, since $\mathbf{v}$ is the concatenation of the vectors on the left, for $i = 1, \ldots, n$, and $\mathbf{w}$ is the concatenation of the vectors on the right, by (Marshall et al., 2011, Prop. 5.A.7)

$$\mathbf{v} \sqsubset \mathbf{w}$$

Hence

$$\begin{aligned}
\sum_j \mathbb{E}\phi^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j &= \sum_j \sum_i w^i_j \mathbb{E}\mathbf{p^i}_j \\
&= \sum_k \mathbf{w}_k \mu_k \\
&\geqslant \sum_k \mathbf{v}_k \mu_k \\
&= \sum_j \sum_i v^i \mathbb{E}\mathbf{p^i}_j = \sum_j \mathbb{E}\lambda^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j
\end{aligned}$$

[29]For vectors $(x_1, \ldots, x_n), (y_1, \ldots, y_n) \in \mathbb{R}^n$, let $x_{[1]} \geqslant \cdots \geqslant x_{[n]}$ and $y_{[1]} \geqslant \cdots \geqslant y_{[n]}$ denote their components in decreasing order. Then $(x_1, \ldots, x_n) \sqsubset (y_1, \ldots, y_n)$ if $\sum_{i=1}^k x_{[i]} \leqslant \sum_{i=1}^k y_{[i]}$ for all $k = 1, \ldots, n$ and $\sum_{i=1}^n x_{[i]} = \sum_{i=1}^n y_{[i]}$, whereas $(x_1, \ldots, x_n) \sqsubset_w (y_1, \ldots, y_n)$ if the first condition holds but not necessarily the second.

where the first equality follows from an argument analogous to that in the proof of Theorem 7, and the inequality follows from Marshall et al. (2011, Prop 4.H.2.c), the previous majorization and the fact that μ is increasing.

Moreover, by Marshall et al. (2011, Prop 5.A.1.b)

$$(\mathbf{v}_1^2, \ldots, \mathbf{v}_{n\times m}^2) \sqsubset_w (\mathbf{w}_1^2, \ldots, \mathbf{w}_{n\times m}^2)$$

whence

$$\begin{aligned}
&\sum_j \mathbb{E}(\phi^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j - \mathbb{E}\phi^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j)^2 \\
&= \sum_j \mathbb{E}(\sum_i w_j^i \mathbf{p^i}_j - \mathbb{E}\sum_i w_j^i \mathbf{p^i}_j)^2 \\
&= \sum_j \sum_i (w_j^i)^2 \mathbb{E}(X_j^i - \mathbb{E}X_j^i)^2 \\
&= \sum_k \mathbf{w}_k^2 \sigma_k \\
&\leqslant \sum_k \mathbf{v}_k^2 \sigma_k \\
&= \sum_j \sum_i (v^i)^2 \mathbb{E}(X_j^i - \mathbb{E}X_j^i)^2 \\
&= \sum_j \mathbb{E}(\sum_i v^i X_j^i - \mathbb{E}\sum_i v^i X_j^i)^2 \\
&= \sum_j \mathbb{E}(\lambda^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j - \mathbb{E}\lambda^{pw}(\mathbf{p^1}, \ldots, \mathbf{p^n})_j)^2
\end{aligned}$$

where the first equality follows from an argument analogous to that in the proof of Theorem 7, the second and second last equalities hold due to Assumption 3, and the inequality follows from Marshall et al. (2011, Prop 4.H.3.b), the previous majorization and the fact that, under Assumption 4, σ is decreasing.

These two observations suffice to establish the result, by the arguments in the proof of Theorem 7.

□

C Proofs of results in Section 6

Proof of Theorem 5. 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. □