

Radical Pooling and Imprecise Probabilities*

Forthcoming at *Erkenntnis*

Ignacio Ojea Quintana

Abstract

This paper focuses on *radical* pooling, or the question of how to aggregate credences when there is a fundamental disagreement about which is the relevant logical space for inquiry. The solution advanced is based on the notion of consensus as common ground (Levi [30]), where agents can find it by suspending judgment on logical possibilities. This is exemplified with cases of scientific revolution. On a formal level, the proposal uses algebraic joins and imprecise probabilities; which is shown to be compatible with the principles of marginalization, rigidity, reverse bayesianism, and minimum divergence commonly endorsed in these contexts. Furthermore, I extend results from previous work by Stewart and Ojea Quintana [43, 44] to show that pooling sets of imprecise probabilities can satisfy important pooling axioms.

1 Introduction

There are various candidate interpretations of an opinion pool in the epistemological literature: a rational consensus; a compromise adopted for the purpose of group decision making; the opinion a group member adopts after learning the opinion of her ‘epistemic peers’; the opinion an external agent adopts upon being informed of the n expert opinions, etc. [18, 47]. The interpretative line endorsed in this essay is the first one: pooling considered as a form of *consensus* for the sake of the argument. Following the literature on pooling, agents’ attitudes are modeled by probability functions over an algebra of events. Much of the literature [9, 18, 47] focused on a constrained form of disagreement among the parties, namely when they disagree *solely* on the degree of belief assigned to events on a *shared* outcome and event space. The purpose of this paper is to generalize the question of pooling to cover cases when agents disagree *radically*; not only on their probabilistic judgments but also on the logical space over which those judgments are made.

Imprecise probabilities allow for an interesting and philosophically well-motivated account of consensus [30, 39]: *consensus as common ground*. Consider first the case of *full* or *plain* belief. Konieczny and Pino-Perez [25] provides a good survey of the literature on belief merging of propositional bases. For each of the propositions in the doxastic space of possibilities, agents are said to have three attitudes: acceptance, rejection, or suspension of judgment. *At the outset of inquiry*, inquirers may seek consensus as *shared agreement* in their beliefs. This could be achieved by retaining whatever beliefs are accepted by all parties and suspending judgment on those that are not shared. Inquiry starting from the consensus view can proceed without begging questions against parties in the consensus, allowing various hypotheses of concern to receive a fair hearing. Such a consensus constitutes a neutral starting point for subsequent inquiry.

*I would like to thank Katie Steele, Michael Nielsen, Rush Stewart, and the anonymous reviewers for their insightful comments. This paper was supported by Australian Research Council Grant DP190101507 and funding from the Humanising Machine Intelligence Grand Challenge Project at the Australian National University.

The idea that parties joining their efforts in inquiry or decision making should restrict themselves to their shared agreements can be extended to judgments of probability in the imprecise probabilities setting. In this setting agents credal attitudes are modeled by sets of probabilities $\{\mathbf{p}_i\}_{i \in I}$. To suspend judgment among some number of probability distributions is to use them for the purposes of inference and decision making. If the parties seeking consensus all agree that a distribution \mathbf{p} is *not* permissible, then the consensus position reflects that agreement and rules it out. A *set* of probability functions represents the shared agreements among the group concerning which probability functions are permissible to use in inference and decision making. This avenue was explored by Stewart and Ojea Quintana [43, 44]. The purpose of this essay is to extend this approach to the case of *radical* pooling.

In this setup, a well informed modeler or arbiter is given a group of agents who disagree radically, with each agent i having an event space, a σ -algebra of propositions $\langle \mathcal{A}_i \rangle$, and a distribution $\langle \mathbf{p}_i \rangle$ on that. Their task is to aggregate or pool the information in those profiles in a way that reflects the consensus as common ground. In the disagreement literature [11], and in the pooling literature [9], it is usually assumed that agents agree on the event space, but not on the probability distribution. I here relax that assumption; agents can be in disagreement about the event space. Intuitively, the event space constitutes the logical space of propositions over which our agents have probabilistic attitudes.¹ Agents are said to disagree *radically* if their event spaces are different. Other forms of disagreement are *non-radical*. The task here is to find a consensus position among parties, a sort of pooling operation over the logical spaces and their credal attitudes towards them. Furthermore, this is a third person perspective project, which is distinct but related to the first person question of *(un)awareness* in epistemology. In cases of (un)awareness a *single* agent grows or contracts their set of possibilities to form a new set of possibilities. These types of belief changes have been explored by Bradley [2], and more recently by Steele & Stefánsson [41] and Mahtani [32] in the probabilistic case; by Cresto [4] in an AGM-style belief revision, and within game theory in the literature by Fagin & Halpern [13], Halpern [20], Dekel *et al* [6], and Modica & Rustichini [34]. Section 3 elaborates on the relations between the project here and issues of awareness growth.

The guiding idea defended here is that the *common ground* consensus involves a sort of suspension of judgment on what is logically possible. Similar to the case of plain belief, the common ground is strongest position that is weak enough to be compatible with everyone's views, and therefore constitutes a good non-question begging starting point for inquiry.

The paper has six sections. The next section motivates the view with an example from the history of science. The third one presents the view of consensus as common ground in detail, acknowledges certain assumptions, justifies the third person perspective, and argues for taking the join of the algebras as the logical pool. Section 4 argues that (i) marginalization, (ii) rigidity and reverse bayesianism, and (iii) divergence lead to the adoption of imprecise probabilities when agents expand their probabilistic judgments to larger algebras. The following section extends results from previous work by Stewart and Ojea Quintana [43, 44] to show that pooling sets of imprecise probabilities can satisfy important pooling axioms. The final section is conclusive.

¹Or for that matter, doxastic attitudes.

2 A Motivating Example

2.1 The Priestley - Lavoisier debate

Finding consensus in the presence of radical disagreement resembles questions about (revolutionary) theory change in philosophy of science. Central to the revolutionary view of science advanced by Kuhn [26, 27] and Feyerabend [14] is the claim that language used in a field of science changes so radically during a revolution that the old language and the new one are not inter-translatable. More generally, they espoused the idea that in those situations opposing theories are *incommensurable*. Similarly, agents disagreeing about the logical space of possibilities over which they make their probabilistic judgments might share no common basis on which to measure their divergence.

The resemblance has its limits. Here pooling is about finding rational consensus between parties. But according to the mainstream reading of Kuhn, in scientific revolutions parties do not have a recourse to a common theoretical language, body of observational evidence, or methodological rules. Despite the differences, the analogy is helpful. I am here concerned with Kuhn’s taxonomic or conceptual incommensurability. During periods of scientific revolution, existing concepts are replaced with new concepts that are at odds with the old ones because they do not share the same “lexical taxonomy”, they cross-classify objects according to different sets of kinds.

In order to illustrate the kind of conceptual tool developed here, I will build on the Priestley-Lavoisier debate around Phlogiston Theory. I will present and formalize the debate in Bayesian terms, and articulate a kind of consensual resolution of the disagreement that does not pretend to be historically accurate of what happened, but that I hope the reader finds sensible.

Priestley advocated the Phlogiston Theory, which attempted to give an account of combustion. Its basic explanatory hypothesis (H) was that there is a substance which is emitted in combustion, namely *phlogiston*. For example, when a metal is heated, phlogiston is emitted, and the calx of the metal is obtained. Furthermore, the process is sometimes reversible. By heating the red calx of mercury, Priestley found that he could obtain the metal mercury, and a new kind of *air* which he called “dephlogisticated air”. In contrast, Lavoisier’s theory of combustion is akin to our modern theory.

Phlogiston Theory (PT)		Modern Theory (MT)	
Input	Output	Input	Output
Metal + air	Calx of metal +	Metal + air	Metal oxide + air
+ heat	phlogisticated air	+ heat	which is poor in oxygen
Red calx (oxide)	Mercury +	Oxide of mercury	Mercury + oxygen
of mercury + heat	dephlogisticated air	+ heat	

It *seems* easy to attain *consensus*. After all, identifying “dephlogisticated air” with “oxygen” (and “calx of metal” with “metal oxide”) seems to provide a good mapping for these experiments. But not all experimental outcomes of dephlogisticated air correspond to experimental outcomes of oxygen, albeit some do. Phlogiston theory departs from the false presupposition (H) that there is a unique substance which is *always* emitted in combustion, and this premise contaminates its terminology and carving of the logical space. This is Kuhn’s and Feyerabend’s point, that terms of phlogiston theory are *theory-laden*: “phlogiston” refers to that which is emitted in all cases of combustion. For Lavoisier, H is false and therefore “phlogiston” fails to refer.

Although they recognize that Phlogiston Theory is incorrect, historians of science might want to explain its success and acknowledge some of its truths. But such a revolutionary theory change from Phlogiston Theory to Modern Theory seems to leave no room for continuity or consensus. Kitcher’s [24] analysis of this case can save us here. Kitcher explains how Priestley’s use of “dephlogisticated air” is theory-laden. Priestley’s early utterances of the expression, driven by the hypothesis H, correspond to “the substance obtained by removing from air the substance which is emitted in combustion”. Those token expressions of “dephlogisticated air” corresponded to a set of referents different from posterior token expressions. *Some* of the later utterances, after his isolation yet misidentification of oxygen through the heating of a red calx (oxide) of mercury, correspond to “the gas obtained by heating the red calx of mercury”. Hence *those* tokens of “dephlogisticated air” refer to oxygen. Lavoisier found a way to interpret the different tokens Priestley used, and therefore carve the space in a more refined way. Furthermore, in so far as Lavoisier *et al.* used the term “oxygen” to refer to those tokens of “dephlogisticated air”, there was room for communication between theoretical rivals.

2.2 Representing the Debate and Finding Consensus

The problem of theory change is being reconceptualized here as one of radical disagreement. Two tasks remain in this introduction: to represent this case in the language of Bayesianism, and to articulate a resolution on the basis of consensus as common ground.

From the perspective of a well informed historian, Priestley and Lavoisier came to the table with different probabilistic attitudes $\mathbf{p}_P, \mathbf{p}_L$ and algebras of propositions $\mathcal{A}_P, \mathcal{A}_L$. In our example, driven by hypothesis H, Priestley can be reconstructed as partitioning his event space so that there is an event P that represents the proposition that the output of the experiment includes phlogisticated air; $P = \{w_i : w_i \text{ s.t. the experiment results in the release of phlogiston into the air}\}$. After heating of a red calx (oxide) of mercury, Priestley is represented as discovering a gas that had a great capacity for absorbing phlogiston, and called it “dephlogisticated air”; $DP = \{w_i : w_i \text{ s.t. the experiment results in the removal of phlogiston from the air}\}$. What matters is that propositions and events like P and DP were at the core of phlogiston theory. As such, both P and DP are taken to be basic, maximally specific, propositions in \mathcal{A}_P .²

The difference, from our modeler’s perspective, is that Lavoisier’s logical carving \mathcal{A}_L was more *refined*. Where Priestley saw a gas without phlogiston, Lavoisier sometimes saw oxygen, *but not always*. Lavoisier understood that “dephlogisticated air” was a misnomer. The proposition DP includes instances of oxygen release but also others. The proposition O that the outcome of the experiment is the release of oxygen entails but it is not entailed by DP , $O \subset DP$. Given H, the absence of phlogiston in the air was at the core of the theory and DP was basic. By rejecting H, Lavoisier could refine the space. In other words, let $DP = \{w_1, w_2\} \in \mathcal{A}_P$. If $\{w_1\}$ corresponds to an instance with oxygen and $\{w_2\}$ without, then $\{w_1\}, \{w_2\} \notin \mathcal{A}_P$. For Lavoisier, $\{w_1\}, \{w_2\} \in \mathcal{A}_L$, so that at least with respect to this issue he made more relevant distinctions.

One initial way of finding consensus would be for our modeler to represent consensus as the more refined logical space \mathcal{A}_L on the grounds that logical refinement constitutes an epistemic improvement. But it is

²A proposition A is basic or atomic in the algebra if it is non-empty and there is no stronger (non-empty) proposition $A' \subset A$ in the algebra. This implies that if both P and DP are basic, their intersection is empty and therefore you can not have the both be true at the same time.

<u>Phlogiston Theory</u>	<u>Modern Theory</u>	<u>Consensus</u>
\mathcal{A}_P	\mathcal{A}_L	$\mathcal{A}_P \vee \mathcal{A}_L$
$DP = \{w_1, w_2\} \in \mathcal{A}_P$	$\{w_1, w_2\} \in \mathcal{A}_L$	$\{w_1, w_2\} \in \mathcal{A}_P \vee \mathcal{A}_L$
$O = \{w_1\} \notin \mathcal{A}_P$	$\{w_1\} \in \mathcal{A}_L$	$\{w_1\} \in \mathcal{A}_P \vee \mathcal{A}_L$

not in general true that a more refined space is an indication of epistemic improvement; including irrelevant distinctions can make a space more convoluted without any epistemic gain. Also, in many interesting circumstances it is not the case that one space is a refinement of the other.

In other words, Priestley may rationally reject Lavoisier’s logical space as the right one on the basis that it is more refined. What would then be a reasonable *consensus*? The adoption of an initial common ground with further investigation in the future. This means retreating to a position that is compatible with all of the individual positions, but that still contains the shared agreement. In the case of algebras, there is a formalism that captures this: the *join*, $\mathcal{A}_L \vee \mathcal{A}_P$, the coarsest algebra that is finer than those two. In detail, its the algebra \mathcal{A}_* such that (i) $\mathcal{A}_L \subseteq \mathcal{A}_*$ and $\mathcal{A}_P \subseteq \mathcal{A}_*$, and (ii) for any \mathcal{A}^* with $\mathcal{A}_L \subseteq \mathcal{A}^*$ and $\mathcal{A}_P \subseteq \mathcal{A}^*$, $\mathcal{A}_* \subseteq \mathcal{A}^*$. Since $\mathcal{A}_P \subset \mathcal{A}_L$, we have that $\mathcal{A}_L \vee \mathcal{A}_P = \mathcal{A}_L$, so Priestley might be convinced to accept Lavoisier’s carving as an initial space, not because it is more refined, but because it constitutes a fair, non question begging, starting point.

The rest of the essay will be devoted to developing and generalizing this notion of consensus for all the relevant cases. Although it is clear that this is not how Lavoisier and Priestley actually resolved their disagreements, it is felicitous that the scientific community adopted \mathcal{A}_L .

3 Finding a Logical Common Ground

3.1 Consensus as Common Ground

For each agent $i \in I$, we, the modelers, consider an *agenda* \mathcal{A}_i , the propositions an agent $i \in I$ regards possible, and a probability function $\mathbf{p}_i : \mathcal{A}_i \rightarrow [0, 1]$ satisfying the Kolmogorov axioms. Propositions in \mathcal{A}_i are the objects of credal and plain belief attitudes. In the plain belief case, consensus as common ground amounts to retaining all of the propositions that are accepted by all of the agents, and suspending judgment on any other proposition. In the probabilistic case, to suspend judgment among some number of distributions is to regard them permissible for the purposes of inference and decision making, and this is captured by a *set* representation, as studied by Stewart and Ojea Quintana [43, 44].

There is an important difference between the case of radical disagreement and that of probabilistic (or for that matter, full belief) disagreement. In the latter, there is agreement on which are the objects of the epistemic attitudes. In other words, there is agreement on which are the relevant propositions of concern, but there might be disagreement about what credal or full belief judgment to have on them. In contrast, in the case of radical disagreement, there is no consensus on what is the relevant set of propositions - i.e. the objects of epistemic attitudes. Or more seriously, there might not be agreement about the *identity* of propositions across logical spaces, following Mahtani [32].

Part of Kitcher’s [24] contribution in the debate about theory change was to point out that (provided we make certain assumptions about the references of terms) we can offer some continuity in cases of scientific revolution [Kuhn [26, 27], Feyerabend [14]]. A well-informed modeler or historian of science might want to explain the success of phlogiston theory, acknowledge some of their truths, and regard its advocates as rational; even in light of the fact that the theory proved to be false. The situation in our case is similar. In order to provide a consensus view we assume that a well-informed modeler can identify propositions across

logical spaces; that some propositions share the same truth-conditions.

This assumption is somewhat warranted by the idea that our arbiter is trying to model rational consensus as common ground. Without assuming *any* identities of possibilities across parties there is not enough room for any form of comparison, much less finding common ground. Arguably, there can be cases like that.³ That does not preclude other forms of consensus to be attainable without that assumption. Consensus can be reached through revolution, conversion, voting, bargaining, or some other psychological, institutional, social, or political process. But that is not finding a common ground, and in many cases it might not be *rational* consensus. It is on the basis of the recognition of identities across spaces that finding common ground is at all possible. This is the main reason the present proposal is that of a third person well-informed modeler trying to find consensus, in contrast first person individuals finding consensus among themselves. Individuals might fail or be unwilling to cross-identify propositions.

Nevertheless, in most of the scenarios we care about, individual agents agree on the meaning of at least some propositions. Suppose two agents share a propositional variable Q_j and agree on its truth conditions. They can begin by agreeing on the same set of possible worlds to assign a reference to it. In our semi-formal example of the Priestley-Lavoisier debate, and in so far as Lavoisier was more refined than Priestley, they then shared the meaning of a significant part of the terminology (but not all). In natural languages this also occurs frequently. Sentences involving general nouns like ‘X is a tenant’ can be recognized as having the same truth-conditions, although different agents might be aware of different ways in which it can be made true. In other words, one agent could point to a way the world might be, which the other agent was not aware of, and they would both agree whether that possibility makes the proposition true or not.

This speaks to the connection between the current third person account of consensus with the first person accounts of awareness. In so far as there is agreement about the meaning of certain propositions between parties, the approach presented here can be proved useful. Furthermore, a well-informed modeler can be viewed as an arbiter or judge mediating disagreement between parties. By first retreating to a position that is non-question begging, this ideal arbiter can later settle the dispute through inquiry. Since the common ground is not at stake, and if the rules of inquiry are agreed upon parties, the disagreement can be in principle resolved. In cases of awareness growth, a single individual becomes aware of new possibilities, and the guiding question is whether there are rational rules for incorporating them to their current logical space. In some cases, but arguably not all, this process can be represented as a disagreement between temporal slices of a single individual. In some other cases, as a disagreement between alternative updates after awareness growth. In addition, a well-informed and neutral (third person) arbiter serves as a regulative idea for a rational first person individual. In the face of disagreement with someone else, or in general confronted with an alternative carving of the logical space, our agent can ask themselves what would an ideal neutral arbiter do. The thought of having an ideal neutral arbiter working as a regulative concept famously appears in Adam Smith’s [40] notion of an *impartial spectator*, although in this case it concerns the appropriateness of actions and behaviors, and not what is the relevant logical carving.

3.2 Joins as Common Ground

Let us assume then that we are given a set of agents each with their respective language \mathcal{L}_i , corresponding to their respective algebras \mathcal{A}_i in a way that some of the terminology shared in the languages has the same truth-conditions. Furthermore, each agent has a subjective probability distribution \mathbf{p}_i on their algebra \mathcal{A}_i , but that will be the subject of the next section. The guiding goal in this section is to provide a *common*

³Cases of true Kuhnian revolution where there are no identities across the parties logical carving.

ground logical space.

The suggestion here is to take the *join* of the \mathcal{A}_i . The *join* of a set of algebras is the coarsest algebra that is finer than all of the individual ones in the set.⁴ By taking the *join* as the shared event space agents are precisely following the heuristic defended here: taking as common ground the coarsest space that is compatible with everyone's original position. Taking any reduced algebra would imply excluding events deemed possible by some agent. Taking any finer algebra would be rationally unjustified.

The analogy with full belief is helpful here. Following the semantic interpretation of Belief Revision Logic AGM, the belief state of an agent can be identified with the (logical closure) of the strongest proposition P in the algebra that is believed by that agent. Given two agents with belief states P_1 and P_2 , finding a common ground means contracting to a state P^* that is weaker than P_1 and P_2 ; this is (i) $P_1 \subseteq P^*$ and $P_2 \subseteq P^*$. But ideally, P^* should also be as strong as possible to preserve the largest common ground, hence for any P_* satisfying (i) it should be the case that $P^* \subseteq P_*$. This is precisely the *join*. In this analogy, having a stand about which propositions are true is akin to having a stand about which event space is the relevant, they are both intentional attitudes about how the world is. Before, a belief state P^* was *stronger* than another belief state P_* just in case $P^* \subseteq P_*$. In our jargon here, an algebra \mathcal{A}_* is stronger than another algebra \mathcal{A}^* just in case $\mathcal{A}_* \subseteq \mathcal{A}^*$. Moving from \mathcal{A}_* to \mathcal{A}^* is an epistemic weakening, since it involves bringing into consideration *more* possibilities as relevant. Finally, finding a common ground is analogous in both cases: fall back to the strongest position compatible with the parties.

Furthermore, the proposed solution aligns with the scientific consensus that resulted from the Priestley-Lavoisier dispute. By rejecting the hypothesis phlogiston is released in combustion, Lavoisier studied the properties of the gas resulting from heating a red calx of (oxide of) mercury as properties of a new substance, oxygen. Priestley studied them as properties of a gas depleted from phlogiston and started with a coarser algebra. But then the join $\mathcal{A}_P \vee \mathcal{A}_L$ is just \mathcal{A}_L and the proposed common ground is precisely Lavoisier's carving. In this we follow Steele & Stefánsson [41] and Mahtani [32], where awareness growth by refinement corresponds to comparing two algebras, one more refined than the other. But comparing one algebra to a more refined one is just one form algebraic disagreement, it is also possible for i to be more refined in some respects, but more coarse in others.⁵

Nothing precludes the alternative that *all* of the parties involved considered possibilities that none of the other did. For example, for two agents A and B we could reasonably have $\Omega^* = \Omega_A = \Omega_B = \{w_1, w_2, w_3, w_4\}$, \mathcal{A}_A is generated by $\{\{w_1\}, \{w_2\}, \{w_3, w_4\}\}$, while \mathcal{A}_B is generated by $\{\{w_1, w_2\}, \{w_3\}, \{w_4\}\}$. Here A distinguishes between $\{w_1\}$ and $\{w_2\}$ while B does not, and conversely for $\{w_3\}$ and $\{w_4\}$. Importantly, the consensus position acknowledges as doxastically possible events that *no one* considered possible before. $\mathcal{A}_A \vee \mathcal{A}_B = \mathcal{A}^* = \mathcal{P}(\Omega^*)$, and in particular $\{w_2, w_3\} \in \mathcal{A}^*$ but $\{w_2, w_3\} \notin \mathcal{A}_A$ and $\{w_2, w_3\} \notin \mathcal{A}_B$.⁶ I take this to be a positive feature of the proposal. It shows that seeking consensus in one of its weakest forms, as common ground and at the outset of inquiry, already can lead to conceptual refinement and innovation.

4 Radical Pooling and Imprecision

Recall that in this setup a well informed modeler or historian of science is given a group of agents who disagree radically (like Priestley and Lavoisier did), with each agent i having an event space $\langle \mathcal{A}_i \rangle$, and a

⁴This is the algebra \mathcal{A}_* such that (i) $\mathcal{A}_L \subseteq \mathcal{A}_*$ and $\mathcal{A}_P \subseteq \mathcal{A}_*$, and (ii) for any \mathcal{A}^* with $\mathcal{A}_L \subseteq \mathcal{A}^*$ and $\mathcal{A}_P \subseteq \mathcal{A}^*$, $\mathcal{A}_* \subseteq \mathcal{A}^*$.


⁵Say $\Omega_1 = \Omega_2 = \{w_1, w_2, w_3, w_4\}$, while $\mathcal{A}_1 = \sigma(\{\emptyset, \{w_1\}, \{w_2\}, \{w_3, w_4\}\})$ and $\mathcal{A}_2 = \sigma(\{\emptyset, \{w_3\}, \{w_4\}, \{w_1, w_2\}\})$.

⁶Notice that although the event $\{w_2, w_3\}$ was not considered possible, some of the states (worlds) constituting the event were considered by the agents. In the language introduced here, this amounts to *partial* radical disagreement, which is different from *full* radical disagreement. The Priestley-Lavoisier debate illustrates the importance of the former.

distribution $\langle p_i \rangle$ on that. The arbiter’s task is to aggregate or pool the information in those profiles in a way that reflects the consensus as common ground. The aggregation procedure in the previous section dealt with the σ -algebra of propositions $\langle \mathcal{A}_i \rangle$ by taking the join. This section will deal with the probabilities $\langle p_i \rangle$ by using imprecise pooling, interpreted here as a form of rational consensus.

4.1 A case for imprecision

The moves so far were the following:



How should the agent extend their subjective probability, originally defined over \mathcal{A}_i to the new common ground logical space \mathcal{A}^* ?

The rest of this section will provide arguments for imprecision. I will show that imprecision is *compatible* with principles of rationality that some argue can be appealed to in these types of cases: marginalization, rigidity, reverse bayesianism, and divergence measures. The present account remains in principle neutral with respect to the issue whether those constraints are rationally required. The purpose is to show how coherent imprecision is with other principles.

The use of imprecise probabilities in the context of disagreement has well known downsides. For example, as Elkin and Wheeler [12] explain, taking convex sets of probability functions may lead to failures of probabilistic independence. Take probability functions p_1 and p_2 such that $p_1(A \cap B) = p_1(A)p_1(B)$ and $p_2(A \cap B) = p_2(A)p_2(B)$. A convex combination $p^* = \alpha p_1 + (1 - \alpha)p_2$ with $0 \leq \alpha \leq 1$, might fail to preserve independence: $p^*(A \cap B) \neq p^*(A)p^*(B)$. A defense against this objection was given in [43]. Furthermore, the source of imprecision here is not due to *pooling* but due to *awareness growth*. Agents will extend their probabilities to posterior algebras that are in general richer than the prior ones. As the discussion on rigidity will reveal, in this case judgments of independence are preserved at this stage.

The proposal endorsed here will be to extend precise probabilities defined on \mathcal{A}_i to sets of probabilities defined on \mathcal{A}^* by taking *natural extensions*. This is an old and familiar move in the Imprecise Pooling literature. This is usually done in the more general language of gambles and previsions. Cooman & Troffaes [5] present a helpful introduction, plus some interesting results on prevision aggregation. Troffaes & Cooman’s book [46] develops extensively the theory of (lower) previsions. It is well known since de Finetti that in general (except degenerate cases) a linear prevision (i.e. a precise probability) on a strict subset of the set of all gambles has no unique natural extension to the larger set: natural extensions are generally imprecise.

The contribution for the rest of the section is to motivate the use of natural extensions on the basis of principles that have been endorsed in the unawareness literature, bridging imprecision and unawareness. Furthermore, natural extensions are usually interpreted as the most conservative correction of the original prevision to a coherent one over the larger set of gambles. This suits well the dialectic endorsed here; when expanding to new algebras agents should be as conservative as they can while still trying to accommodate for the new information.

4.2 Marginalization leads to imprecision

Given three σ -algebras $\mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \mathcal{A}_3$ a probability function \mathbf{p} defined over \mathcal{A}_2 , the marginalization of \mathbf{p} on \mathcal{A}_1 ($\mathbf{p} \upharpoonright_{\mathcal{A}_1}$) is the unique \mathbf{p}_* defined over \mathcal{A}_1 such that for all elements $E \in \mathcal{A}_1$, $\mathbf{p}(E) = \mathbf{p}_*(E)$.⁷ Conversely, we can look at all of the probabilities \mathbf{p}^* defined on \mathcal{A}_3 such that restricted or marginalized to \mathcal{A}_2 is equal to \mathbf{p} , i.e. $\mathbf{p}^* \upharpoonright_{\mathcal{A}_2} = \mathbf{p}$. That is what we do here.

In our setup, each probability measure \mathbf{p}_i is defined over \mathcal{A}_i , and the task is to extend that measure to \mathcal{A}^* . We will use imprecise probabilities. In particular, there are several admissible extensions of \mathbf{p}_i to \mathcal{A}^* and there is no reason for i to prefer one over the other. In our notation, \mathbf{p}_i^* refers to distributions defined over \mathcal{A}^* , the richer algebra.

$$\mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\}$$

So \mathbb{P}_i is the set of all probability measures that extend \mathbf{p}_i - i.e. all the \mathbf{p}_i^* defined over \mathcal{A}^* such that restricted to \mathcal{A}_i gives back \mathbf{p}_i . More precisely: $\mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i$ if and only if for each $H \in \mathcal{A}_i$, $\mathbf{p}_i^*(H) = \mathbf{p}_i(H)$.⁸

Proposition 1. \mathbb{P}_i is convex for all i .

New events will have non-null probability weight only if they have old events as subsets. Completely new events will have null probabilistic weight. In the imagined scenario, Priestley would assign zero probability to Lavoisier's alleged observation.

4.3 Rigidity (or Reverse Bayesianism) and Conservatism lead to Imprecision

A standard, yet contested, principle of Bayesian epistemology is *rigidity*: conditional probabilities are kept fixed throughout updating. Let $\mathbf{p}(\cdot)$ be defined over some algebra \mathcal{A} of events be a representation of some agent's (coherent) credal state. Suppose now the agent learns an event E and updates their credal state to $\mathbf{q}(\cdot) = \mathbf{p}(\cdot|E)$ according to Bayesian standards. For all events $H \in \mathcal{A}$ we have:

- **Rigidity:** $\mathbf{q}(H|E) = \mathbf{p}(H|E \cap E) = \mathbf{p}(H|E)$

Richard Jeffrey [21] famously argued against this form of rigidity. I do not intend to discuss this form of rigidity, but one relevant for the radical disagreement case.

Intuitions are not clear on whether rigidity is an acceptable rationality constraints in the context of awareness growth. Nevertheless, there are defenses of rigidity-like principles in the literature, in particular in Bradley [2] and Steele and Stefánsson [41]. As explained before, although the approach defended here sympathizes with rigidity, the central purpose is to show that it is compatible with it.

The idea behind this principle is that we should extend our relational attitudes to the new set in such a way as to conserve all prior relational credences. Becoming aware of the new events should not affect the relative credibility of previously considered events. For all $E, H \in \mathcal{A}_i$ with $\mathbf{p}_i(E) \neq 0 \neq \mathbf{p}_i^*(E)$:

- **Strong Radical Rigidity [SRR]:** $\mathbf{p}_i^*(H|E) = \mathbf{p}_i(H|E)$

SRR can arguably be regarded as too strong of a requirement. If so, consider the principle that demands that the strengthened conditional probabilities of old events given the old set of possibilities should be kept the same. Namely, for all $H \in \mathcal{A}_i$:

⁷ \mathbf{p}_* need not be a probability function.

⁸In general the marginalization operation $\mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i}$ is defined only when \mathcal{A}_i is a subalgebra of \mathcal{A}^* , but this needs not to be the case here and that is why the definition is necessary.

- **Weak Radical Rigidity [WRR]:** $\mathbf{p}_i^*(H|\Omega_i) = \mathbf{p}_i(H|\Omega_i) = \mathbf{p}_i(H)$

The following observation holds trivially, provided that $\mathbf{p}_i^*(\Omega_i) \neq 0$.

Proposition 2. *Strong Radical Rigidity and Weak Radical Rigidity are equivalent.*

Related with rigidity is the principle of *Reverse Bayesianism* introduced by Karni and Vierø [22, 23]. More precisely, they extend the Savage framework by introducing extra axioms regulating awareness and they show that the collection of those axioms admit an expected utility representation in which the probability at hand obeys *reverse bayesianism*. A natural way to translate such property from the Savage framework to the present is the following:

- **Reverse Bayesianism [RB]:** For all $A, B \in \mathcal{A}_i$ non-null with respect to \mathbf{p}_i and \mathbf{p}_i^* :

$$\frac{\mathbf{p}_i(A)}{\mathbf{p}_i(B)} = \frac{\mathbf{p}_i^*(A)}{\mathbf{p}_i^*(B)}$$

The property states that the likelihood ratios of events in the original space remains intact after the expansion to the new space. Importantly, Karni and Vierø [22, 23] results show that the probabilistic representation has the property for *non-null* [non zero probability] events in *both* \mathbf{p}_i and \mathbf{p}_i^* and for *both* A and B (not just the denominator).

Proposition 3. *Strong Radical Rigidity implies Reverse Bayesianism.*

Proposition 4. *Reverse Bayesianism does not imply Strong Radical Rigidity in general.*

Proposition 5. *Reverse Bayesianism implies Strong Radical Rigidity, for all $E, H, E \cap H \in \mathcal{A}_i$ non-null with respect to \mathbf{p}_i and \mathbf{p}_i^* .*

Proposition 6. *Unconstrained Reverse Bayesianism⁹ is equivalent to Strong Radical Rigidity.*

Hence for our purposes here, accounts compatible with rigidity will also be compatible with Reverse Bayesianism.

Several authors, in particular Steele and Stefánsson [42, 41] argue that Reverse Bayesianism is not an acceptable general principle for awareness change. In particular, they offer counter examples in which the principle fails with common sense intuitions. These objections are fair for RB as a *general* principle, but are not particularly harmful to the argument made here. First, I am not arguing that RB is a good *general* principle for radical revisions, but a reasonable one in the context of opinion pooling. Second, as with marginalization, the argument is meant to show that the use of imprecise probabilities is *compatible* with RB [SRR, WRR].¹⁰

Going back to imprecision, it is easy to observe that all the $\mathbf{p}_i^* \in \mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\}$ satisfy SRR,¹¹ yet the set of all \mathbf{p}_i^* satisfying SRR is strictly larger than \mathbb{P}_i .¹² The following condition suffices to make the two sets coextentional.

⁹This means generalizing RB so that only the denominator is required to be non-null.

¹⁰Yet, if *only* marginalization, RB, SRR, WRR and divergence measures are required, then imprecision will be required too.

¹¹If $\mathbf{p}_i^* \in \mathbb{P}_i$, then for all $H \in \mathcal{A}_i$, $\mathbf{p}_i^*(H) = \mathbf{p}_i(H)$. Hence, $\mathbf{p}_i^*(H|E) = \frac{\mathbf{p}_i^*(H \cap E)}{\mathbf{p}_i^*(E)} = \frac{\mathbf{p}_i(H \cap E)}{\mathbf{p}_i(E)} = \mathbf{p}_i(H|E)$.

¹²Let $\Omega_i = \{w_1, w_2\}$, $\mathcal{A}_i = \mathcal{P}(\Omega_i)$, and $\mathbf{p}_i(w_1) = \frac{1}{2}$. On the other hand, let $\Omega^* = \{w_1, w_2, w_3\}$, $\mathcal{A}^* = \mathcal{P}(\Omega^*)$, and $\mathbf{p}_i^*(w_1) = \mathbf{p}_i^*(w_2) = \mathbf{p}_i^*(w_3) = \frac{1}{3}$. Then \mathbf{p}_i^* satisfies SRR but $\mathbf{p}_i^* \notin \mathbb{P}_i$.

- **Radical Conservatism:** $\mathbf{p}_i^*(\Omega_i) = \mathbf{p}_i(\Omega_i) = 1$.

Proposition 7. *The set of all \mathbf{p}_i^* defined over \mathcal{A}^* satisfying SRR [equivalently, unconstrained RB] and Radical Conservatism is just $\mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\}$.*

Radical Conservatism requires that outcomes originally inaccessible to the agent will get null probability, which requires some justification in this context.

First, this is *not* yet the output of a *pooling* function. Our modeler is not taking into account the *credences* of the agents but only their algebras. No (probabilistically representable) evidence in support of the new possibilities is being considered, and *probabilistic* consensus is not being sought. This will be done at the next stage.

Second, the guiding notion of consensus is one in which agents minimally weaken their credal positions so as to accommodate those of other agents. But an agent i did not originally acknowledge all the $w \in \Omega^*$ but $w \notin \Omega_i$ as relevant doxastic possibilities, so there seems to be no reason to give them any probabilistic weight. Also, as pointed before, giving them positive probabilistic weight implies reducing the probabilistic weight of doxastic possibilities they originally considered relevant; which is not a conservative operation.

4.4 Divergence Measures lead to Imprecision

Another way of understanding the question on how to extend \mathbf{p}_i from \mathcal{A}_i to \mathcal{A}^* is by taking \mathbf{p}_i to be a *partial* credence function over \mathcal{A}^* , one in which there are no values assigned for some elements of the algebra. A partial function can be considered *incoherent* if it assigns values to some elements but not their unions, intersections or complements.¹³ Hence the question of how to structurally strengthen \mathbf{p}_i to \mathcal{A}^* is akin to the issue of how to fix incoherent credences, a topic extensively developed by Pettigrew [37] and Predd *et al* [38]. Their solution involves using divergence measures.

The argument here is not that using divergence is a rational *requirement*, but that imprecision is compatible with it. In most cases divergence measures, or more generally different forms of loss functions, are warranted on the basis that we are trying to better estimate some value like truth (in the *accuracy first* literature) or some target function (in function approximation in supervised and unsupervised learning). The problem here is not of estimation but rather of *coherence* [37] and conservatism: among the coherent distributions over the richer algebra, which is the set of closest to the original distribution over the prior algebra.

In this section, \mathbb{P} is the set of *credence* functions defined on \mathcal{A} (rather than \mathcal{A}^*).¹⁴ A divergence is a function $\mathcal{D}_{\mathcal{A}} : \mathbb{P} \times \mathbb{P} \rightarrow [0, \infty]$ such that for all credences $\mathbf{p}, \mathbf{p}' \in \mathbb{P}$, (i) $\mathcal{D}_{\mathcal{A}}(\mathbf{p}, \mathbf{p}) = 0$, and (ii) $\mathcal{D}_{\mathcal{A}}(\mathbf{p}, \mathbf{p}') > 0$ if $\mathbf{p} \neq \mathbf{p}'$. In general, neither symmetry nor the satisfaction of triangle inequality are required. The following are two classical examples:

Squared Euclidean Distance [SED]

$$SED_{\mathcal{A}}(\mathbf{p}, \mathbf{p}') = \sum_{E \in \mathcal{A}} [\mathbf{p}(E) - \mathbf{p}'(E)]^2$$

¹³This is related yet different to coherence as sure loss avoidance.

¹⁴Notice that nothing in the definition requires the elements of \mathbb{P} to be *probability* functions, just that they assign values in $[0,1]$ to the elements of \mathcal{A}

Generalized Kullback-Leibler [GKL]

$$GKL_{\mathcal{A}}(\mathbf{p}, \mathbf{p}') = \sum_{E \in \mathcal{A}} \left[\mathbf{p}(E) \log \left(\frac{\mathbf{p}(E)}{\mathbf{p}'(E)} \right) - \mathbf{p}(E) + \mathbf{p}'(E) \right]$$

Divergence measures are generally used under the methodological assumption of minimal mutilation, which states that any shift from a prior to a posterior credal states should accommodate the posterior condition by minimally changing the prior. For example, Bayesian conditionalization satisfies the minimal mutilation principle according to GKL: $\mathbf{p}(\cdot|E) = \mathbf{p}'$ minimizes $GKL_{\mathcal{A}}(\mathbf{p}, \mathbf{p}')$ on the condition that the \mathbf{p}' is a probability function (i.e. coherent credal functions) that assigns value 1 to E . Similarly, Kullback-Leibler divergence generalizes Jeffrey conditionalization [7]. GKL is also a generalization of Jaynes' Maximum Entropy formalism [48].

To “fix” \mathbf{p}_i amounts then to finding the (set of) coherent \mathbf{p}_i^* defined over \mathcal{A}^* that minimizes some appropriate divergence measure. We can avoid the question on whether to use SED or GKL by looking at a generalization of both:

Additive Bregman Divergence [ABD]

Suppose $\phi : [0, 1] \rightarrow \mathbb{R}$ is a strictly convex function that is twice differentiable on $(0,1)$ with a continuous second derivative. Suppose $\mathcal{D} : \mathbb{P} \times \mathbb{P} \rightarrow [0, \infty]$. Then \mathcal{D} is the additive Bregman divergence generated by ϕ if, for any $\mathbf{p}, \mathbf{p}' \in \mathbb{P}$,

$$\mathcal{D}_{\mathcal{A}}(\mathbf{p}, \mathbf{p}') = \sum_{E \in \mathcal{A}} [\phi(\mathbf{p}(E)) - \phi(\mathbf{p}'(E)) - \phi'(\mathbf{p}'(E))(\mathbf{p}(E) - \mathbf{p}'(E))]$$

SED is the ABD generated by $\phi(x) = x^2$, and GKL is the ABD generated by $\phi(x) = x \log x - x$. The additive Bregman divergence can also be justified on accuracy-first grounds. Predd *et al* [38] show that if \mathbf{p} is an incoherent credence function, then the coherent credence function \mathbf{p}' that minimizes ABD with respect to it is more accurate than \mathbf{p} at all possible worlds. Therefore “fixing” incoherent credence functions using ABD increases accuracy. The following observation is important for our purposes:

Proposition 8. *For all \mathbf{p}_i^* coherent credence functions defined over \mathcal{A}^* :*

$$\mathbf{p}_i^* \in \mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\} \text{ iff } \mathcal{D}_{\mathcal{A}_i}(\mathbf{p}_i, \mathbf{p}_i^*) = 0 = \mathcal{D}_{\mathcal{A}_i}(\mathbf{p}_i^*, \mathbf{p}_i)$$

This section explored the structural strengthening of \mathbf{p}_i to \mathcal{A}^* by extending it to $\mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\}$. This was shown to be *compatible* with marginalization, the principle of minimal mutilation expressed by divergence measures, and some conservative constrains involving rigidity and reverse bayesianism. Any of the $\mathbf{p}_i^* \in \mathbb{P}_i$ will satisfy those principles. Furthermore, excluding any of the \mathbf{p}_i^* satisfying the conditions would be arbitrary and unjustified. If we take any or all of those conditions to be exhaustive principles for structural strengthening, then we would be required to adopt $\mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i\}$.

4.5 Extending imprecise sets of probability functions

The previous sections focused on extending probability measures \mathbf{p}_i defined over \mathcal{A}_i to \mathcal{A}^* . But the reader might endorse the view according to which subjective individual credal states should be represented by sets

of probability functions - in particular *convex* sets. So suppose we start with a convex set of probabilities P_i representing agent i 's credences over \mathcal{A}_i and we would like to extend it to \mathcal{A}^* . The proposed extension is:

$$\mathbb{P}_i = \{\mathbf{p}_i^* : \mathbf{p}_i^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i \text{ for some } \mathbf{p}_i \in P_i\}$$

The following proposition is helpful:


Proposition 9. *Provided that P_i is convex, \mathbb{P}_i is convex.*

5 Imprecise Pooling

So far, each agent has a convex set of probability functions \mathbb{P}_i defined over a common algebra \mathcal{A}^* . Say $\mathcal{P}(P)$ is the power set of the set of all probability functions defined on an algebra \mathcal{A}^* . A generalized imprecise pooling function is defined as:

$$\mathcal{F} : (\mathcal{P}(P))^n \rightarrow \mathcal{P}(P)$$

Our goal in this section is to find the appropriate pooling function \mathcal{F} to conclude the aggregation process:



The next section provides an introduction to pooling, enumerating some well know results and principles adopted in the literature. Importantly, much of the literature assumes that the input of the pooling function is a precise profile of probabilities. We here have imprecise inputs. The sections following will develop the account for doing this, concluding with some generalized results.

5.1 Pooling Basics

Say P is the set of probability functions defined over a common algebra \mathcal{A} . Formally, a *precise* pooling method for a group of n individuals is a function:

$$F : P^n \rightarrow P$$

mapping profiles of probability functions for the n agents $(\mathbf{p}_1, \dots, \mathbf{p}_n)$, to *single* probability functions $F(\mathbf{p}_1, \dots, \mathbf{p}_n)$. The probabilities are assigned to events, represented by a common algebra \mathcal{A} defined over a shared sample space Ω .

Various concrete pooling functions have been studied. Linear pooling functions may be the most common and obvious proposal [45, 33, 28].

Linear Opinion Pools. $F(\mathbf{p}_1, \dots, \mathbf{p}_n) = \sum_{i=1}^n w_i \mathbf{p}_i$, where $w_i \geq 0$ and $\sum_{i=1}^n w_i = 1$.

Another proposal is to take a weighted *geometric* instead of a weighted arithmetic average of the n probability functions [31, 1, 17].¹⁵

Geometric Opinion Pools. $F(\mathbf{p}_1, \dots, \mathbf{p}_n) = c \prod_{i=1}^n \mathbf{p}_i^{w_i}$, where $w_i \geq 0$ and $\sum_{i=1}^n w_i = 1$, and $c = \frac{1}{\sum_{\omega' \in \Omega} [\mathbf{p}_1(\omega')]^{w_1} \dots [\mathbf{p}_n(\omega')]^{w_n}}$ is a normalization factor.¹⁶

A third, more recent proposal from Dietrich [8] is given by the following formula.

Multiplicative Opinion Pools. $F(\mathbf{p}_1, \dots, \mathbf{p}_n)(\omega) = c \prod_{i=1}^n \mathbf{p}_i$, where \mathbf{p}_0 is a fixed ‘‘calibrating’’ probability function, and $c = \frac{1}{\sum_{\omega' \in \Omega} [\mathbf{p}_0(\omega)] \cdot [\mathbf{p}_1(\omega')] \dots [\mathbf{p}_n(\omega')]}$ is a normalization factor.¹⁷

On the other hand, several principles have been suggested:

Strong Setwise Function Property (SSFP): There exists a function $G : [0, 1]^n \rightarrow [0, 1]$ such that, for every event A , $F(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = G(\mathbf{p}_1(A), \dots, \mathbf{p}_n(A))$.

Weak Setwise Function Property (WSFP): There exists a function $G : \mathcal{A} \times [0, 1]^n \rightarrow [0, 1]$ such that, for every event A , $F(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = G(A, \mathbf{p}_1(A), \dots, \mathbf{p}_n(A))$.

Zero Preservation Property (ZPP): For any event A , if $\mathbf{p}_i(A) = 0$ for $i = 1, \dots, n$, then $F(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = 0$.

Marginalization Property (MP): Let \mathcal{A}' be a sub- σ -algebra of \mathcal{A} . For any $A \in \mathcal{A}'$, $F(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = F([\mathbf{p}_1 \upharpoonright_{\mathcal{A}'}], \dots, [\mathbf{p}_n \upharpoonright_{\mathcal{A}'}])(A)$.

Probabilistic Irrelevance Preservation (IP): If $\mathbf{p}_i(A|B) = \mathbf{p}_i(A)$ for $i = 1, \dots, n$, then $F^B(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = F(\mathbf{p}_1, \dots, \mathbf{p}_n)(A)$.¹⁸

Unanimity Preservation (UP): For all $(\mathbb{P}_1, \dots, \mathbb{P}_n) \in \mathbb{P}^n$, if $\mathbb{P}_i = \mathbb{P}_j$ for all $i, j = 1, \dots, n$, then $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = \mathbb{P}_i$.

External Bayesianity (EB): For every profile $(\mathbf{p}_1, \dots, \mathbf{p}_n)$ in the domain of F and every likelihood function λ such that $(\mathbf{p}^\lambda, \dots, \mathbf{p}_n^\lambda)$ remains in the domain of F , $F(\mathbf{p}_1^\lambda, \dots, \mathbf{p}_n^\lambda) = F^\lambda(\mathbf{p}_1, \dots, \mathbf{p}_n)$.¹⁹

There is a large literature justifying these principles, and studying the interaction between them and the different forms of precise pooling. Surveys of the literature include Simon French’s [15], Genest and Zidek’s [18], and Dietrich and List’s [9]. This is not the place to cover that material. The reason to invoke the principles is that there are limitative results showing that the axioms can not be jointly satisfied by precise pooling functions.

Theorem 1. [33, Theorem 3.3; 28, Theorem 6.7] *Given that the algebra contains at least three disjoint events, a pooling function satisfies SSFP iff it is a linear pooling function.*

¹⁵An unweighted geometric pool of n numerical values is given by $\sqrt[n]{x_1 \dots x_n} = x_1^{\frac{1}{n}} \dots x_n^{\frac{1}{n}}$.

¹⁶The domain of geometric pooling operators is restricted to profiles of *regular* pmfs, i.e., those \mathbf{p} such that $\mathbf{p}(\omega) > 0$ for all $\omega \in \Omega$. We denote the set of regular pmfs P' making the relevant domain P'^n .

¹⁷As with geometric pooling functions, the domain of multiplicative pooling functions will be restricted to P'^n .

¹⁸Here $F^B(\mathbf{p}_1, \dots, \mathbf{p}_n)$ refers to updating the pool function $F(\mathbf{p}_1, \dots, \mathbf{p}_n)$ with the event B .

¹⁹The requirement is that updating the individual probabilities on a common *likelihood function* (as opposed to updating on an event) λ and then pooling is the same as pooling and then updating the pool on that likelihood function. A likelihood function, $\lambda : \Omega \rightarrow [0, \infty)$, is intended to encode, given any $\omega \in \Omega$, how expected some evidence is with the number $\lambda(\omega)$. In conditionalizing, $\lambda(\cdot)$ serves the same role as the conditional probability $\mathbf{p}(E|\cdot)$ in Bayes’ theorem.

Theorem 2. [33, Corollary 3.4] *Given that the algebra contains at least three disjoint events, F satisfies WSFP and ZPP iff F is a linear pooling function.*

McConway has shown that a pooling function has the WSFP iff it has the MP. So linear pooling functions satisfy MP and ZPP.

Theorem 3. [16, p. 1104] *The geometric pooling functions are externally Bayesian and preserve unanimity.*

Part of the appeal of using imprecise probabilities, as shown by Stewart and Ojea Quintana [43], is that they can jointly satisfy (an appropriate translation of) criteria that precise pooling cannot:

Theorem 4. [43, Propositions 2,3,5]

Convex IP pooling functions satisfy SWFP, WSFP, MP, ZPP, and unanimity preservation.

Convex IP pooling functions are externally Bayesian.

*Convex IP pooling functions satisfy irrelevance preservation.*²⁰

The same strategy will be used in Section 5.3, when considering how to aggregate *sets* of precise probabilities.

5.2 Pooling Imprecise Probabilities

We can now move to our problem at hand. Now we are given a profile of *sets* of probabilities functions \mathbb{P}_i defined over a common algebra \mathcal{A}^* . This problem has been explored in the literature, but not particularly in the current setting.

In an unpublished research note [36], Walley provides an account of information fusion and presents fifteen criteria for the aggregation of opinions represented as sets of probabilities. The report is very hard to get, but both Dubois *et al.* [10] and Seamus Bradley [3] provide a summary of it, together with general setup for information fusion. None of them present the problem as one of finding consensus, and they are also not focused on finding it in the context of radical disagreement. Furthermore, they are not in direct dialogue with the literature on pooling and disagreement presented in the previous section. Nevertheless, and maybe unsurprisingly, some of the criteria coincide with those in pooling; we will get into that in short. Walley [36] comes up with a number of merging rules that result from the fifteen criteria. Some of them, the most tractable ones in the pooling setting, are:

Proposal 1: $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = \bigcap_{i=1}^n \mathbb{P}_i$.

Proposal 2 $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = \bigcup_{i=1}^n \mathbb{P}_i$.

Proposal 3 $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = H(\bigcup_{i=1}^n \mathbb{P}_i)$.²¹

Moral and del Sagrado [35] also consider postulates for merging credal sets inspired by Walley. Their proposal is too involved to elaborate here, but two contrasting points are in place. First, their setup is one of expert opinion merging, and they adopt Dempster-Shafer theory of evidence in order to weight each agent according to their expertise. The approach here is one in which agents are all taken into account fairly and equally, and not given different weights. The second point is that they consider postulates that will fail here. For example, they postulate that if \mathbb{P}^* is the set of all probabilities over \mathcal{A}^* , then $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n, \mathbb{P}^*) = \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$. The idea is that \mathbb{P}^* represents complete ignorance and therefore should be ignored in merging. Walley [36] and Dubois *et al.* [10] consider versions of this principle. Proposals 2 and 3 clearly fail at this.

²⁰This particular principles does require considering probabilistic *irrelevance* rather than probabilistic *independence*.

²¹ $H(A)$ here is the convex closure of A .

In the present interpretation of the operator as one of consensus, even an ignorant (or possibly skeptical) agent with credal state \mathbb{P}^* should be given a fair hearing. Taking unions (Proposal 2) ensures this. Taking the convex set of unions (Proposal 3) might lead to greater imprecision, but this might be warranted if all potential resolutions are to be preserved at the outset. The next few sections will be a defense of Proposals 2 and 3 on the basis of pooling axioms.

5.3 The Pooling Axioms Reformulated

As mentioned before, Walley [36] offered several criteria for aggregating imprecise probabilities, as do Moral and del Sagrado [35] and Bradley [3]. We are here going to follow the pooling framework, and reformulate the axioms presented in Section 5.1 for aggregating sets of probabilities:

Strong Setwise Function Property (SSFP): There exists a function $G : \mathcal{P}([0, 1])^n \rightarrow \mathcal{P}([0, 1])$ such that for any $A \in \mathcal{A}^*$, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = G(\mathbb{P}_1(A), \dots, \mathbb{P}_n(A))$.

Weak Setwise Function Property (WSFP): There exists a function $G : \mathcal{A}^* \times \mathcal{P}([0, 1])^n \rightarrow \mathcal{P}([0, 1])$ such that, for any $A \in \mathcal{A}^*$, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = G(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A))$.

Zero Preservation Property (ZPP): For any $A \in \mathcal{A}^*$, if $\mathbb{P}_i(A) = \{0\}$ for all $i = 1, \dots, n$, then $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \{0\}$.

Marginalization Property (MP): Let \mathcal{A}'^* be a sub- σ -algebra of \mathcal{A}^* . For any $A \in \mathcal{A}'^*$, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{F}([\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}], \dots, [\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}])(A)$.²²

Irrelevance Preservation (IP): If $\mathbf{p}(A|B) = \mathbf{p}(A)$ for all $\mathbf{p} \in \bigcup_{i=1}^n \mathbb{P}_i$, then $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{F}^B(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$.

Unanimity Preservation (UP): For all $(\mathbb{P}_1, \dots, \mathbb{P}_n) \in \mathbb{P}^n$, if $\mathbb{P}_i = \mathbb{P}_j$ for all $i, j = 1, \dots, n$, then $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = \mathbb{P}_i$.

External Bayesianity (EB): For every profile $(\mathbb{P}_1, \dots, \mathbb{P}_n)$ in the domain of \mathcal{F} and every likelihood function λ such that $(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda)$ remains in the domain of \mathcal{F} , $\mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda) = \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$.²³

Some of the principles considered by Walley [36] coincide, most notably commutativity with updating (External Bayesianity here). Some other are not explicit but still hold for our cases, like Symmetry (the order of the input does not affect the output). And some others are in Walley but not in pooling or vice versa. The argument for Proposals 2 and 3 will be that they both satisfy all of the pooling axioms reformulated.

5.4 Some Results

Most of these results are extensions of the ones present in Stewart and Ojea Quintana [43].

Proposition 10. *Let $\mathcal{F} : (\mathcal{P}(P))^n \rightarrow \mathcal{P}(P)$ be an IP pooling function (not necessarily convex). \mathcal{F} satisfies WSFP iff \mathcal{F} satisfies MP.*

Proposition 11. *Convex IP pooling functions (Proposal 3) satisfy SSFP, WSFP, MP, ZPP, and unanimity preservation.*

²²Here $[\mathbb{P}_i \upharpoonright_{\mathcal{A}'^*}] = \{\{\mathbf{p}_i \upharpoonright_{\mathcal{A}'^*}\} : \mathbf{p}_i \in \mathbb{P}_i\}$.

²³Here $\mathbb{P}_i^\lambda = \{\mathbf{p}_i^\lambda : \mathbf{p}_i \in \mathbb{P}_i\}$.

Lemma 1. [Cf. 29, 19] Convexity is preserved under updating on a likelihood function, i.e., $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is convex.

Proposition 12. Convex IP pooling functions (Proposal 3) are externally Bayesian.

Proposition 13. Convex IP pooling functions (Proposal 3) satisfy irrelevance preservation.


Proposition 14. Let $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = \bigcup_{i=1}^n \mathbb{P}_i$ (Proposal 2). Then, \mathcal{F} satisfies SSFP, WSFP, ZPP, MP, unanimity preservation, external Bayesianity, and confirmational irrelevance preservation.²⁴

It is trivial to observe that Proposal 1 fails to satisfy some of the axioms.

The question of how to aggregate sets of probabilities in general is still open in the literature. There are open problems about the interpretation of the operation and about which are the adequate axioms or principles for it. The strategy in this section was to defend Proposals 2 and 3 on the basis that they satisfy reasonable criteria for pooling, and an interpretation of pooling in terms of consensus as common ground. Taking unions and convex hulls of unions represents taking the strongest position compatible with all of the individual ones.

6 Conclusion

The general procedure can be visualized in the following diagram:



The nature of this essay is programmatic. It is focused in assessing a form of disagreement that has been neglected in the literature, disagreement about which is the relevant logical space of possibilities. The position defended here is that pooling can be interpreted as a technical and philosophical characterization of a form of consensus, namely consensus as common ground at the outset of inquiry.

This view was exemplified with the Priestley and Lavoisier debate in Section 2. Section 3 presents the view of consensus as common ground in detail, acknowledges certain assumptions, justifies the third person perspective, and argues for taking the join of the algebras as the logical pool. Section 4 argues that (i) marginalization, (ii) rigidity and reverse bayesianism, and (iii) divergence lead to the adoption of imprecise probabilities when agents expand their probabilistic judgments to larger algebras. The previous section extends results from previous work by Stewart and Ojea Quintana [43, 44] to show that pooling sets of imprecise probabilities can satisfy important pooling axioms.

²⁴The proof of this proposition is straightforward and so is omitted here. Also, Proposal 2 satisfies a stronger version of irrelevance preservation, *stochastic* preservation: If $\mathbf{p}(A \cap B) = \mathbf{p}(A)\mathbf{p}(B)$, for all $\mathbf{p} \in \bigcup_{i=1}^n \mathbb{P}_i$, then, for all $\mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$, $\mathbf{p}(A \cap B) = \mathbf{p}(A)\mathbf{p}(B)$.

As explained in Section 3, the proposal assumed the perspective of a third person well-informed modeler that was capable of identifying propositions and their truth conditions across different logical spaces. As discussed there, this is a necessary conditions in order to attain consensus of the form we are interested here. Assumptions of this form, discovered by Mahtani [32], are still being discussed in the literature and may have further philosophical consequences.

It is important to remark that taking consensus as common ground is not the only solution to the problem, but it is a sound one. On the one hand, there might be other formal approaches. For example, Cooman & Troffaes [5] consider different (possibly incompatible) sources of information about independent variables and propose an aggregation method considering product algebras and probabilistic (lower previsions) extensions to them. On the other hand, there might be other philosophical approaches, by giving another interpretation to the aggregation process. One might think, for example, that an agent that has a prior defined over a more refined event or sample space has some kind of epistemic priority over a coarser (less “sophisticated”) agent. Conversely, an expert’s algebra may be coarser because it removed irrelevant distinctions.

The Priestley and Lavoisier debate served as a motivating example of the kind of disagreements that can be found in science, or inquiry in general. I argued that some cases of conceptual incommensurability can be modeled as a disagreement about what are the relevant sample and event spaces, what are the best natural joints to carve. But even in the face of such discrepancy, some form of consensus can be defined. At the outset, disagreements about what variables are relevant for explanation should begin by giving a fair treatment to all the possibilities. During inquiry, some of those possibilities may be discarded. But a pooling function cannot be expected to give an armchair account of the outcome of inquiry. In order to resolve radical disagreement rationally agents need start the conversation by taking others’ world views as possible.

References

- [1] Michael Bacharach. Scientific disagreement. *Unpublished Manuscript*, 1972.
- [2] Richard Bradley. *Decision Theory with a Human Face*. Cambridge University Press, 2017.
- [3] Seamus Bradley. Aggregating belief models. In *Proceedings of ISIPTA 2019*, Proceedings of Machine Learning Research, 2019.
- [4] Eleonora Cresto. A model for structural changes of belief. *Studia Logica*, 88(3):431–451, 2008.
- [5] Gert De Cooman and Matthias Troffaes. Coherent lower previsions in systems modelling: Products and aggregation rules. *Reliability Engineering & System Safety*, 85:113–134, 07 2004.
- [6] Eddie Dekel, Barton Lipman, and Aldo Rustichini. Standard state-space models preclude unawareness. *Econometrica*, 66(1):159–174, 1998.
- [7] Persi Diaconis and Sandy L. Zabell. Updating subjective probability. *Journal of the American Statistical Association*, 77(380):822–830, 1982.
- [8] Franz Dietrich. Bayesian group belief. *Social Choice and Welfare*, 35(4):595–626, 2010.
- [9] Franz Dietrich and Christian List. Probabilistic Opinion Pooling. In Alan Hájek and Christopher Hitchcock, editors, *Oxford Handbook of Probability and Philosophy*. Oxford University Press, 2014.
- [10] Didier Dubois, Weiru Liu, Jianbing Ma, and Henri Prade. The basic principles of uncertain information fusion. an organised review of merging rules in different representation frameworks. *Information Fusion*, 32:12 – 39, 2016.
- [11] Adam Elga. Reflection and Disagreement. *Noûs*, 41(3):478–502, 2007.
- [12] Lee Elkin and Gregory Wheeler. Resolving peer disagreements through imprecise probabilities. *Noûs*, 2018.
- [13] Ronald Fagin and Joseph Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39 – 76, 1987.
- [14] Paul K. Feyerabend. Explanation, reduction and empiricism. In H. Feigl and G. Maxwell, editors, *Crítica: Revista Hispanoamericana de Filosofía*, pages 103–106. 1962.
- [15] Simon French. Group Consensus Probability Distributions: A Critical Survey. In M.H. DeGroot D.V. Lindley J.M. Bernardo and A.F.M. Smith, editors, *Bayesian Statistics: Proceedings of the Second Valencia International Meeting*, volume 2, pages 183–201. North-Holland, September 1985.
- [16] Christian Genest. A Characterization Theorem for Externally Bayesian Groups. *The Annals of Statistics*, pages 1100–1105, 1984.
- [17] Christian Genest, Kevin J McConway, and Mark J Schervish. Characterization of Externally Bayesian Pooling Operators. *The Annals of Statistics*, pages 487–501, 1986.
- [18] Christian Genest and James V Zidek. Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, pages 114–135, 1986.

- [19] Francisco Javier Girón and Sixto Ríos. Quasi-Bayesian Behaviour: A More Realistic Approach to Decision Making? *Trabajos de Estadística y de Investigación Operativa*, 31(1):17–38, 1980.
- [20] Joseph Y. Halpern. Alternative semantics for unawareness. *Games and Economic Behavior*, 37(2):321 – 339, 2001.
- [21] R.C. Jeffrey. *The Logic of Decision*. McGraw-Hill series in probability and statistics. University of Chicago Press, 1990.
- [22] Edi Karni, Quitzé Valenzuela-Stookey, and Marie-Louise Viero. Reverse bayesianism: A generalization. 12 2018.
- [23] Edi Karni and Marie-Louise Viero. "reverse bayesianism": A choice-based theory of growing awareness. *American Economic Review*, 103, 01 2011.
- [24] Philip Kitcher. Theories, theorists and theoretical change. *The Philosophical Review*, 87(4):519–547, 1978.
- [25] Sébastien Konieczny and Ramón Pino Pérez. Logic based merging. *Journal of Philosophical Logic*, 40:239–270, 04 2011.
- [26] Thomas S. Kuhn. *The structure of scientific revolutions*. University of Chicago Press, Chicago, 1970.
- [27] Thomas S. Kuhn. *The Road since Structure*. University of Chicago Press, 2000.
- [28] Keith Lehrer and Carl Wagner. *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*, volume 21. Springer, 1981.
- [29] Isaac Levi. Irrelevance. In *Foundations and Applications of Decision Theory*, pages 263–273. Springer, 1978.
- [30] Isaac Levi. Consensus as Shared Agreement and Outcome of Inquiry. *Synthese*, 62(1):pp. 3–11, 1985.
- [31] Albert Madansky. Externally Bayesian groups. 1964.
- [32] Anna Mahtani. Awareness growth and dispositional attitudes. *Synthese*, pages 1–17, forthcoming.
- [33] Kevin J McConway. Marginalization and Linear Opinion Pools. *Journal of the American Statistical Association*, 76(374):410–414, 1981.
- [34] Salvatore Modica and Aldo Rustichini. Unawareness and partitional information structures. *Games and Economic Behavior*, 27(2):265 – 298, 1999.
- [35] Serafín Moral and José del Sagrado. *Aggregation of Imprecise Probabilities*, pages 162–188. Physica-Verlag HD, Heidelberg, 1998.
- [36] Walley P. The elicitation and aggregation of beliefs. *Technical Report*, 1982.
- [37] Richard Pettigrew. Aggregating incoherent agents who disagree. *Synthese*, Nov 2017.
- [38] Joel B. Predd, Robert Seiringer, Elliott H. Lieb, Daniel N. Osherson, H. Vincent Poor, and Sanjeev R. Kulkarni. Probabilistic coherence and proper scoring rules. *IEEE Trans. Inf. Theor.*, 55(10):4786–4792, October 2009.

- [39] Teddy Seidenfeld, Joseph B Kadane, and Mark J Schervish. On the shared preferences of two Bayesian decision makers. *The Journal of Philosophy*, 86(5):225–244, 1989.
- [40] A. Smith and K. Haakonssen. *Adam Smith: The Theory of Moral Sentiments*. Cambridge Texts in the History of Philosophy. Cambridge University Press, 2002.
- [41] Katie Steele and H. Orri Stefánsson. Belief revision for growing awareness. *Mind*, forthcoming.
- [42] Katie Steele and Orri Stefansson. *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge University Press.
- [43] Rush T. Stewart and Ignacio Ojea Quintana. Probabilistic Opinion Pooling with Imprecise Probabilities. *Journal of Philosophical Logic*, page Forthcoming, 2016.
- [44] Rush T. Stewart and Ignacio Ojea Quintana. Learning and Pooling, Pooling and Learning. *Unpublished Manuscript*, MS2.
- [45] Mervyn Stone. The opinion pool. *The Annals of Mathematical Statistics*, 32(4):1339–1342, 1961.
- [46] Matthias CM Troffaes and Gert De Cooman. *Lower previsions*. Wiley series in probability and statistics. John Wiley and Sons, 2014.
- [47] C Wagner. Jeffrey conditioning and external bayesianity. *Logic Journal of IGPL*, 18(2):336–345, 2009.
- [48] P. M. Williams. Bayesian conditionalisation and the principle of minimum information. *The British Journal for the Philosophy of Science*, 31(2):131–144, 1980.

A Proofs

Proof. Proposition 1

Suppose $\mathbf{p}^1, \mathbf{p}^2 \in \mathbb{P}_i$. I want to show that for $\alpha \in [0, 1]$, $\mathbf{p}^* = \alpha \cdot \mathbf{p}^1 + (1 - \alpha) \cdot \mathbf{p}^2 \in \mathbb{P}_i$.

It is enough to show that $\mathbf{p}^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i$. Given any $E \in \mathcal{A}_i$, $\mathbf{p}^*(E) = \alpha \cdot \mathbf{p}^1(E) + (1 - \alpha) \cdot \mathbf{p}^2(E) = \alpha \cdot \mathbf{p}_i(E) + (1 - \alpha) \cdot \mathbf{p}_i(E) = \mathbf{p}_i(E)$. □

Proof. Proposition 2

SRR entails WRR trivially, since $\Omega_i \in \mathcal{A}_i$ and $\mathbf{p}_i(\Omega_i) = 1$. WRR and the definition of conditional probabilities secure that for all $E, H \in \mathcal{A}_i$, $\mathbf{p}_i(H|E) = \frac{\mathbf{p}_i(H \cap E)}{\mathbf{p}_i(E)} = \frac{\mathbf{p}_i(H \cap E|\Omega_i)}{\mathbf{p}_i(E|\Omega_i)} = \frac{\mathbf{p}_i^*(H \cap E|\Omega_i)}{\mathbf{p}_i^*(E|\Omega_i)} = \frac{\mathbf{p}_i^*(H \cap E \cap \Omega_i)/\mathbf{p}_i^*(\Omega_i)}{\mathbf{p}_i^*(E \cap \Omega_i)/\mathbf{p}_i^*(\Omega_i)} = \frac{\mathbf{p}_i^*(H \cap E \cap \Omega_i)}{\mathbf{p}_i^*(E \cap \Omega_i)} = \frac{\mathbf{p}_i^*(H \cap E)}{\mathbf{p}_i^*(E)} = \mathbf{p}_i^*(H|E)$. □

Proof. Proposition 3

Since both A and B are non-null with respect to \mathbf{p}_i and \mathbf{p}_i^* , $\mathbf{p}_i(A|B), \mathbf{p}_i(B|A), \mathbf{p}_i^*(A|B)$ and $\mathbf{p}_i^*(B|A)$ are all well defined. Furtherm Suppose that SRR is true so that $\mathbf{p}_i(A|B) = \mathbf{p}_i^*(A|B)$ and $\mathbf{p}_i(B|A) = \mathbf{p}_i^*(B|A)$.

SRR also ensures that $\mathbf{p}(A \cap B) \neq 0 \neq \mathbf{p}^*(A \cap B)$ or $\mathbf{p}(A \cap B) = 0 = \mathbf{p}^*(A \cap B)$.

Suppose $\mathbf{p}(A \cap B) \neq 0 \neq \mathbf{p}^*(A \cap B)$, then $\mathbf{p}_i(A|B), \mathbf{p}_i(B|A), \mathbf{p}_i^*(A|B)$ and $\mathbf{p}_i^*(B|A)$ are all non-null. The argument is then: $\mathbf{p}_i(B \cap A) \cdot \mathbf{p}_i^*(A \cap B) = \mathbf{p}_i(A \cap B) \cdot \mathbf{p}_i^*(B \cap A)$ iff (by definition of conditional probability) $(\mathbf{p}_i(A) \cdot \mathbf{p}_i(B|A)) \cdot (\mathbf{p}_i^*(B) \cdot \mathbf{p}_i^*(A|B)) = (\mathbf{p}_i(B) \cdot \mathbf{p}_i(A|B)) \cdot (\mathbf{p}_i^*(A) \cdot \mathbf{p}_i^*(B|A))$ iff (using the identities provided by SRR and non-nullity) $\mathbf{p}_i(A) \cdot \mathbf{p}_i^*(B) = \mathbf{p}_i(B) \cdot \mathbf{p}_i^*(A)$ iff $\frac{\mathbf{p}_i(A)}{\mathbf{p}_i(B)} = \frac{\mathbf{p}_i^*(A)}{\mathbf{p}_i^*(B)}$.

Now assume $\mathbf{p}(A \cap B) = 0 = \mathbf{p}^*(A \cap B)$. Classically, $\mathbf{p}(A \cup B) = \mathbf{p}(A) + \mathbf{p}(B) - \mathbf{p}(A \cap B)$. Using the assumption we get $\mathbf{p}(A \cup B) = \mathbf{p}(A) + \mathbf{p}(B)$ and similarly $\mathbf{p}^*(A \cup B) = \mathbf{p}^*(A) + \mathbf{p}^*(B)$. By SRR $\mathbf{p}(A|A \cup B) = \mathbf{p}^*(A|A \cup B)$. So $\frac{\mathbf{p}(A)}{\mathbf{p}(A) + \mathbf{p}(B)} = \frac{\mathbf{p}^*(A)}{\mathbf{p}^*(A) + \mathbf{p}^*(B)}$. Hence, after some basic algebra, $\mathbf{p}_i(A) \cdot \mathbf{p}_i^*(B) = \mathbf{p}_i^*(A) \cdot \mathbf{p}_i(B)$, which is all we need. □

Proof. Proposition 4

Assume RB holds for $A, B \in \mathcal{A}_i$ non-null with respect to \mathbf{p}_i and \mathbf{p}_i^* [i.e. $\mathbf{p}_i(A) \neq 0 \neq \mathbf{p}_i(B)$ and $\mathbf{p}_i^*(A) \neq 0 \neq \mathbf{p}_i^*(B)$], but that there is H with $H \cap E$ null in \mathbf{p}_i but not in \mathbf{p}_i^* (or viceversa). Then $\mathbf{p}_i^*(H|E) \neq \mathbf{p}_i(H|E)$ and SRR fails. □

Proof. Proposition 5

This follows easily by taking $A = E \cap H$ and $B = E$. Notice that $A = E \cap H$ needs to be non-null on both credences for this argument to work. □

Proof. Proposition 6

Suppose we generalize RB so that only the denominator is required to be non-null.

On the one hand, taking $A = E \cap H$ and $B = E$ would do the trick to show RB implies SRR.

On the other hand, assume SRR and suppose for contradiction that unconstrained RB does not hold for $A, B \in \mathcal{A}$: $\frac{\mathbf{p}_i(A)}{\mathbf{p}_i(B)} \neq \frac{\mathbf{p}_i^*(A)}{\mathbf{p}_i^*(B)}$. Clearly, it can't be the case that $\mathbf{p}(A) = 0 = \mathbf{p}^*(A)$. Alternatively, consider the case where $\mathbf{p}(A) = 0$ but $\mathbf{p}^*(A) \neq 0$. Using WRR [equivalent to SRR] we have that $0 = \mathbf{p}(A) = \mathbf{p}(A|\Omega) = \mathbf{p}^*(A|\Omega) = \frac{\mathbf{p}^*(A)}{\mathbf{p}^*(\Omega)} \neq 0$. So we again get a contradiction, and the same would hold in the case where $\mathbf{p}(A) \neq 0$ but $\mathbf{p}^*(A) = 0$. The fourth and final case in the one in which $\mathbf{p}(A) \neq 0 \neq \mathbf{p}^*(A)$; but this is just constrained RB which was shown to be implied by SRR in Proposition 4. □

Proof. Proposition 7

Note 12 showed that the $\mathbf{p}_i^* \in \mathbb{P}_i$ satisfy SRR, and they trivially satisfy Structural Conservatism. Suppose \mathbf{p}_i^* satisfies SRR and Structural Conservatism. Let $H \in \mathcal{A}_i$ (so $H \subseteq \Omega_i$) and consider $\mathbf{p}_i^*(H)$. By Structural Conservatism $\mathbf{p}_i^*(H) = \mathbf{p}_i^*(H|\Omega_i)$. By SRR and Observation 4, \mathbf{p}_i^* satisfies WRR and therefore $\mathbf{p}_i^*(H|\Omega_i) = \mathbf{p}_i(H)$. □

Proof. Proposition 8

The observation follows from the fact that ABD is a divergence measure as defined before. If $\mathbf{p}_i^* \in \mathbb{P}_i$ then for all $E \in \mathcal{A}_i$, $\mathbf{p}_i^*(E) = \mathbf{p}_i(E)$, and hence each element of the summation will be null. On the other hand, if $\mathbf{p}_i^* \notin \mathbb{P}_i$, then there is $E \in \mathcal{A}_i$ such that $\mathbf{p}_i^*(E) \neq \mathbf{p}_i(E)$. I will rely on the well known fact that ϕ is strictly convex if and only if $\forall x \neq y \in [0, 1], \phi(x) > \phi(y) + \phi'(y)(x - y)$. This implies that for any $E \in \mathcal{A}_i$, if $\mathbf{p}_i^*(E) \neq \mathbf{p}_i(E)$, then $\phi(\mathbf{p}(E)) - \phi(\mathbf{p}^*(E)) - \phi'(\mathbf{p}^*(E))(\mathbf{p}(E) - \mathbf{p}^*(E)) > 0$. \square

Proof. Proposition 9

Suppose $\mathbf{p}^1, \mathbf{p}^2 \in \mathbb{P}_i$. I want to show that for $\alpha \in [0, 1]$, $\mathbf{p}^* = \alpha \cdot \mathbf{p}^1 + (1 - \alpha) \cdot \mathbf{p}^2 \in \mathbb{P}_i$.

By the supposition we get that there are $\mathbf{p}_i^1, \mathbf{p}_i^2 \in P$ with $\mathbf{p}_i^1(A) = \mathbf{p}^1(A)$ and $\mathbf{p}_i^2(A) = \mathbf{p}^2(A)$ for all $A \in \mathcal{A}_i$. Since P is convex, $\mathbf{p}_i^* = \alpha \cdot \mathbf{p}_i^1 + (1 - \alpha) \cdot \mathbf{p}_i^2 \in P$.

But then for any $A \in \mathcal{A}_i$, $\mathbf{p}^*(A) = \alpha \cdot \mathbf{p}^1(A) + (1 - \alpha) \cdot \mathbf{p}^2(A) = \alpha \cdot \mathbf{p}_i^1(A) + (1 - \alpha) \cdot \mathbf{p}_i^2(A) = \mathbf{p}_i^*(A)$. So $\mathbf{p}^* \upharpoonright_{\mathcal{A}_i} = \mathbf{p}_i^*$ and $\mathbf{p}^* \in \mathbb{P}_i$. \square

Proof. Proposition 10

This is an adaptation McConway's [33] proof.

WSFP \Rightarrow MP. Assume that \mathcal{F} has the WSFP, i.e., there is a function $G : \mathcal{A}^* \times \mathcal{P}([0, 1]^n) \rightarrow \mathcal{P}([0, 1])$ such that $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = G(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A))$. By WSFP, we have $\mathcal{F}([\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}], \dots$

, $[\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}])(A) = G(A, [\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}](A), \dots, [\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}](A))$. Since G is a function and $\mathbb{P}_i(A) = [\mathbb{P}_i \upharpoonright_{\mathcal{A}'^*}](A)$ for any $A \in \mathcal{A}'^*$ (all such $A \in \mathcal{A}'^*$ are also in \mathcal{A}^*), it follows that $G(A, [\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}](A), \dots, [\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}](A)) = G(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A)) = \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$. Hence, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{F}([\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}], \dots, [\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}])(A)$.

MP \Rightarrow WSFP. Assume that \mathcal{F} has the MP. Let $A \in \mathcal{A}^*$. We want to show that $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$ depends only on A and $\mathbb{P}_i(A), i = 1, \dots, n$.

First, if $A = \emptyset$ or $A = \Omega$, then, since the range of \mathcal{F} is $\mathcal{P}(\mathbb{P})$, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$ depends only on A and $\mathbb{P}_i(A), i = 1, \dots, n$, for any profile because, setting $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{G}(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A))$ and $\mathcal{F}(\mathbb{P}'_1, \dots, \mathbb{P}'_n)(A) = \mathcal{G}(A, \mathbb{P}'_1(A), \dots, \mathbb{P}'_n(A))$, it follows that $\mathcal{G}(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A)) = \mathcal{G}(A, \mathbb{P}'_1(A), \dots, \mathbb{P}'_n(A))$. [i.e. $\mathcal{G}(A, \mathbb{P}_1(A), \dots, \mathbb{P}_n(A)) = \mathcal{G}(A, \mathbb{P}'_1(A), \dots, \mathbb{P}'_n(A))$ is just $\{0\}$ or $\{1\}$.]

Next, suppose that $\emptyset \neq A \neq \Omega$. Consider the σ -algebra $\mathcal{A}'^* = \{\emptyset, A, A^c, \Omega\}$. \mathcal{A}^* contains A and has \mathcal{A}'^* as a sub-algebra. By MP, then

$$\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{F}([\mathbb{P}_1 \upharpoonright_{\mathcal{A}'^*}], \dots, [\mathbb{P}_n \upharpoonright_{\mathcal{A}'^*}])(A).$$

\mathcal{A}'^* is uniquely defined by A and any probability over \mathcal{A}' is uniquely determined by the probability of A under that distribution. So the righthand side of the equation above is determined by A and $[\mathbb{P}_i \upharpoonright_{\mathcal{A}'^*}](A) = \mathbb{P}_i(A)$. \square

Proof. Proposition 11

Recall $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \{\mathbf{p}(A) : \mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)\}$. We let G of the SSFP be $G(\mathbb{P}_1(A), \dots, \mathbb{P}_n(A)) = H(\bigcup_{i=1}^n \mathbb{P}_i(A))$, the convex hull operation applied over $\bigcup_{i=1}^n \mathbb{P}_i(A)$. It is clear that G depends just on the probabilities $\mathbb{P}_i(A), i = 1, \dots, n$. We need to show that

$$\{\mathbf{p}(A) : \mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)\} = H\left(\bigcup_{i=1}^n \mathbb{P}_i(A)\right).$$

But $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = H(\bigcup_{i=1}^n (\mathbb{P}_i))$ by assumption. So the identity is trivial.

But since SSFP clearly implies WSFP, WSFP is satisfied, too. By Proposition 10, it follows immediately that \mathcal{F} has the MP.

Because $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is a set of probability functions, $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(\emptyset) = \{0\}$. Let $\mathbb{P}_i(A) = 0, i = 1, \dots, n$. Since there is a function, \mathcal{G} , such that $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{G}(\mathbb{P}_1(A), \dots, \mathbb{P}_n(A))$, we have it that

$$\begin{aligned}
\{0\} &= \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(\emptyset) \\
&= \mathcal{G}(\mathbb{P}_1(\emptyset), \dots, \mathbb{P}_n(\emptyset)) \\
&= \mathcal{G}(\mathbb{P}_1(A), \dots, \mathbb{P}_n(A)) \\
&= \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)
\end{aligned}$$

So, ZPP follows from SSFP.

For any profile $(\mathbb{P}_1, \dots, \mathbb{P}_n) \in \mathbb{P}^n$, if all \mathbb{P}_i are identical, then the convex hull is just $\{\mathbb{P}_i\}$. So \mathcal{F} satisfies *unanimity preservation*. \square

Proof. Lemma 1

This generalizes a proof due originally to Girón and Rios and Levi [29, 19].

We want to show that given any two members, $\mathbf{p}^\lambda, \mathbf{p}'^\lambda \in \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$ and $\alpha \in [0, 1]$, $\mathbf{p}^\star = \alpha\mathbf{p}^\lambda + (1-\alpha)\mathbf{p}'^\lambda$ is in $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$.

If there is a convex combination of \mathbf{p} and \mathbf{p}' , $\mathbf{p}_\star = \beta\mathbf{p} + (1-\beta)\mathbf{p}'$, such that $\mathbf{p}_\star^\lambda = \mathbf{p}^\star$, then the convexity of $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is established as a consequence of the convexity of $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$.

$$\text{Where } \mathbf{p}_\star^\lambda(\cdot) = \frac{\mathbf{p}_\star(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \mathbf{p}_\star(\omega')\lambda(\omega')} = \frac{\beta\mathbf{p}(\cdot)\lambda(\cdot) + (1-\beta)\mathbf{p}'(\cdot)\lambda(\cdot)}{\beta\sum_{\omega' \in \Omega} \mathbf{p}(\omega')\lambda(\omega') + (1-\beta)\sum_{\omega' \in \Omega} \mathbf{p}'(\omega')\lambda(\omega')}.$$

For any α we want to find some β such that:

$$\mathbf{p}^\star(\cdot) = \alpha\mathbf{p}^\lambda(\cdot) + (1-\alpha)\mathbf{p}'^\lambda(\cdot) = \frac{\beta\mathbf{p}(\cdot)\lambda(\cdot) + (1-\beta)\mathbf{p}'(\cdot)\lambda(\cdot)}{\beta\sum_{\omega' \in \Omega} \mathbf{p}(\omega')\lambda(\omega') + (1-\beta)\sum_{\omega' \in \Omega} \mathbf{p}'(\omega')\lambda(\omega')} = \mathbf{p}_\star^\lambda(\cdot).$$

For $\beta = \frac{\alpha\sum_{\omega^* \in \Omega} \mathbf{p}'(\omega^*)\lambda(\omega^*)}{\alpha\sum_{\omega^* \in \Omega} \mathbf{p}'(\omega^*)\lambda(\omega^*) + (1-\alpha)\sum_{\omega^* \in \Omega} \mathbf{p}(\omega^*)\lambda(\omega^*)}$, the equality is verifiable with some tedious algebra. \square

Proof. Proposition 12

We want $\mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda) = \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$.

$\mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda) \subseteq \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$. It is clear that for each $i = 1, \dots, n$, $\mathbb{P}_i^\lambda \subseteq \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$, since $\mathbb{P}_i \subseteq \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$. By Lemma 1, $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is convex. It follows that $\mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda) \subseteq \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$.

We now want $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n) \subseteq \mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda)$.

$\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) = H(\bigcup_{i=1}^n \mathbb{P}_i)$. Hence every $\mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is a convex combination of elements in $\bigcup_{i=1}^n \mathbb{P}_i$. Let us denote $\mathbf{p} = \sum_j \alpha_j \mathbf{p}_j$, with $\mathbf{p}_j \in \bigcup_{i=1}^n \mathbb{P}_i$.

On the other hand, $\mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n) = \{\mathbf{p}^\lambda : \mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n) \text{ and } \mathbf{p}^\lambda(\cdot) = \frac{\mathbf{p}(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \mathbf{p}(\omega')\lambda(\omega')}\}$.

We show that any $\mathbf{p} \in \mathcal{F}^\lambda(\mathbb{P}_1, \dots, \mathbb{P}_n)$ is identical to some member of $\mathcal{F}(\mathbb{P}_1^\lambda, \dots, \mathbb{P}_n^\lambda)$.

$$\begin{aligned}
\mathbf{p}^\lambda(\omega) &= \frac{\mathbf{p}(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \mathbf{p}(\omega')\lambda(\omega')} && \text{[Definition]} \\
&= \frac{\sum_j \alpha_j \mathbf{p}_j(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \sum_j \alpha_j \mathbf{p}_j(\omega')\lambda(\omega')} \\
&= \frac{\sum_j \alpha_j \mathbf{p}_j^\lambda(\omega) \cdot \sum_{\omega' \in \Omega} \mathbf{p}_j(\omega')\lambda(\omega')}{\sum_{\omega' \in \Omega} \sum_j \alpha_j \mathbf{p}_j(\omega')\lambda(\omega')} && [\mathbf{p}_j(\omega)\lambda(\omega) = \mathbf{p}_j(\omega)^\lambda \cdot \sum_{\omega' \in \Omega} \mathbf{p}_j(\omega')\lambda(\omega')] \\
&= \sum_j \beta_j \mathbf{p}_j^\lambda(\omega) \in \mathcal{F}(\mathbf{p}_1^\lambda, \dots, \mathbf{p}_n^\lambda) && \text{[Since each } \mathbf{p}_j \in \bigcup_{i=1}^n \mathbb{P}_i, \mathbf{p}_j^\lambda \in \bigcup_{i=1}^n \mathbb{P}_i^\lambda\text{]}
\end{aligned}$$

where $\beta_j = \frac{\alpha_j \cdot \sum_{\omega' \in \Omega} \mathbf{p}_j(\omega')\lambda(\omega')}{\sum_{\omega' \in \Omega} \sum_j \alpha_j \mathbf{p}_j(\omega')\lambda(\omega')}$ with $\beta_j \geq 0$ for all $j = 1, \dots, n$ and $\sum_j \beta_j = 1$. \square

Proof. Proposition 13

Notice first that *Irrelevance* preservation is not *Stochastic* preservation. Stochastic preservation states that if $\mathbf{p}(A \cap B) = \mathbf{p}(A)\mathbf{p}(B)$, for all $\mathbf{p} \in \bigcup_{i=1}^n \mathbb{P}_i$, then, for all $\mathbf{p} \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)$, $\mathbf{p}(A \cap B) = \mathbf{p}(A)\mathbf{p}(B)$.

Suppose that $\mathbf{p}(A|B) = \mathbf{p}(A)$ for all $\mathbf{p} \in \bigcup_{i=1}^n \mathbb{P}_i$. We want $\mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A) = \mathcal{F}^B(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$. Consider $\mathbf{p}^*(A) \in \mathcal{F}(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$ and $\mathbf{p}_*(A) \in \mathcal{F}^B(\mathbb{P}_1, \dots, \mathbb{P}_n)(A)$. By convexity we have that $\mathbf{p}^*(A) = \sum_j \alpha_j \mathbf{p}_j(A)$ with all the $\mathbf{p}_j \in \bigcup_{i=1}^n \mathbb{P}_i$, for appropriate α_j .

By Proposition 12 (external Bayesianity), $\mathcal{F}^B(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = \mathcal{F}(\mathbf{p}_1^B, \dots, \mathbf{p}_n^B)(A)$. So $\mathbf{p}_*(A) \in \mathcal{F}(\mathbf{p}_1^B, \dots, \mathbf{p}_n^B)(A)$. Again using convexity we have that $\mathbf{p}_*(A) = \sum_j \beta_j \mathbf{p}_j^B(A)$, with $\mathbf{p}_j^B \in \bigcup_{i=1}^n \mathbb{P}_i^B$ and for appropriate β_j .

By hypothesis $\mathbf{p}_i^B(A) = \mathbf{p}_j(A)$ for all $\mathbf{p}_j \in \bigcup_{i=1}^n \mathbb{P}_i$, $\mathbf{p}_j^B \in \bigcup_{i=1}^n \mathbb{P}_i^B$. Hence, $\mathbf{p}_*(A) = \sum_j \beta_j \mathbf{p}_j^B(A) = \sum_j \beta_j \mathbf{p}_j(A)$. Letting $\alpha_j = \beta_j$, it follows that $\mathcal{F}(\mathbf{p}_1, \dots, \mathbf{p}_n)(A) = \mathcal{F}^B(\mathbf{p}_1, \dots, \mathbf{p}_n)(A)$. \square