Essays on Aggregation in Deliberation and Inquiry

Rush T. Stewart

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in the Graduate School of Arts and Sciences.

COLUMBIA UNIVERSITY

2017

Contents

List of Appendices

Acknowledgments

I would like to express my gratitude, first of all, to my committee members—John Collins, Philip Kitcher, Isaac Levi, Hans Rott, and Achille Varzi—not only for their care and attention in reviewing this work, but also for their guidance. Perhaps without even realizing it, John provided valuable encouragement at several points during my time at Columbia. He graciously agreed to co-advise my work with Isaac, and tolerated a late change to the topic of this dissertation. My interest in certain social aspects of epistemology was inspired in part by some of Philip's ideas. I am very grateful to Philip for the various ways in which he has helped me since I was his TA for Science and Religion. The first time that I served as teaching assistant it was for Achille. I wish to thank him for his advice on a number of occasions over the years since. Hans's work is a model of formal philosophy for its combination of formal rigor and philosophical seriousness. I have learned a tremendous amount from it. I thank him for hosting me in Regensburg in the spring of 2014.

I owe a special debt of gratitude to Isaac. I regard Isaac Levi as one of the most serious, sophisticated, and important philosophers of the last century. In many ways, I was concerned to extend and explore many of his insights in this dissertation. I'm extremely grateful for the opportunity to work so closely with one of my philosophical heroes. My regular meetings with Isaac at his apartment are some of the fondest experiences of my life. I am very fortunate to have two luminaries of pragmatism as members of my dissertation committee.

Certain others, while not on my committee, took no less of an interest in my philosophical education and development. Jeff Helzner was extremely generous to me during our overlapping time at Columbia. I learned a lot about how to pick projects and structure research from him. And I thank him for encouraging me to explore both pragmatism and rational choice early on. I met Greg Wheeler in Lisbon in the summer of 2011. He, perhaps as much as anyone else, has been helping to guide my research and professional development ever since. I'm very thankful to both Jeff and Greg for their help and their examples of integrity.

I have learned so much from friends, colleagues, and mentors. Yang Liu and I founded a reading group in formal philosophy in the fall of 2010, and that group developed into perhaps the most significant source of learning in my graduate education. Thanks to Haim Gaifman for coordinating the Formal Formal Philosophy Group, and thanks to all of the (Informal) Formal Philosophy Group—Robby Finley, Arthur Heller, Yang Liu, Michael Nielsen, and Ignacio Ojea Quintana—for inspiring and instructive conversations. Thanks in particular to Michael and Ignacio for collaborating with me on a number of papers. Some of my joint papers with them have been published (a version of Chapter 2 in the *Journal of Philosophical Logic* and a version of Chapter 3 in *Erkentnnis*), and others are already in press. I am thankful that I learned early in my career how fruitful and enjoyable collaborative work can be, even if it remains less popular in philosophy than in other disciplines. I have no doubt that my work with them is much stronger than I could have made it on my own. Thanks also to other fellow graduate students at Columbia for conversations, advice, and friendship, not least Jorge Morales (my "favorite") and Porter Williams. Matt Duncan, Paul Pedersen, Teddy Seidenfeld, Reuben Stern, and Mark Swails all generously read, commented on, and talked with me about my work at several points. And thanks to all of my coaches and training partners in grappling for the healthy distraction.

Finally, I'd like to express my deepest gratitude to my family. My sister, Meghan Denney, and my mother, Tammy Baker, provided a nearly unreasonable amount of support and understanding during my time in graduate school. I thank my wife Nicolle Faver for her love and companionship and return them.

To my father

CHAPTER 1

Introduction

Mathematical aggregation frameworks are general and precise settings in which to study ways of forming a *consensus* or group point of view from a set of potentially diverse points of view. Aggregation is an important topic in economics (e.g., allocation aggregation), political science (e.g., social choice theory), and statistics (e.g., opinion pooling), but it also finds application throughout decision theory and epistemology. Consensus between points of view is central to any developed account of rationality. While it may be obvious that common ground or consensus matters in settings of group decision-making and non-question-begging joint inquiry, such notions are also crucial for single agents. A rational agent can have multiple goals or values that come into conflict or she may suspend judgment between various ways of evaluating events with respect to subjective probability, for example.

Yet the standard aggregation frameworks have significant limitations. A number of results show that certain sets of desirable aggregation properties cannot be simultaneously satisfied. In this dissertation, I propose general aggregation frameworks, and I investigate their philosophical and mathematical properties. Drawing on work in the theory of *imprecise probabilities*, I make a case for employing *sets* of probability functions and *sets* of preference relations for identifying common ground among candidate ways of evaluating events and options, respectively. Among other things, I prove that many of the famous "impossibility" results afflicting standard models of aggregation (like Arrow's Impossibility Theorem or Sen's Impossibility of the Paretian Liberal in preference aggregation and similar limitative results in probability pooling) do not apply. In other words, certain very attractive constraints on aggregation are simultaneously satisfiable in the generalized aggregation frameworks that I develop.

This dissertation is a sustained exploration of taking indeterminacy or, more generally, suspension of judgment—in probability, full beliefs, and values—seriously in the context of aggregation. It comprises six essays, naturally grouped into three sections, one section on each of probability, full belief, and preference. The first section concerns probability aggregation or *pooling* and contains three essays. Chapter 2 philosophically motivates the use of imprecise probabilities in the context of pooling and generalizes the canonical mathematical framework to allow for set-valued pooling functions. A number of simple possibility results are established and a distinguished format of pooling with imprecise probabilities is characterized.

In Chapter 3, I study commutativity of pooling and learning for the general account of pooling on offer and various probabilistic learning rules. Positive results are stated for Bayesian conditionalization (and a mild generalization of it), imaging, and a certain parameterization of Jeffrey conditioning. This last observation is obtained by slightly generalizing Wagner's characterization of (precise) externally Bayesian pooling operators to the setting of imprecise probabilities. These results might be taken to strengthen the case that pooling should go by imprecise probabilities since no precise pooling method is as versatile.

Chapter 4 can be seen as addressing a potential objection to the use of imprecise probabilities in the context of aggregation. One might complain that pooling with imprecise probabilities allows very *weak* consensus positions, too weak, perhaps, to immediately address the inference or decision problems at hand. One line of response to this objection appeals to asymptotic results that have played important roles in Bayesian theory. These "merging of opinions" theorems show that, for sufficiently similar priors, initial differences of opinions "wash out" in the long run when conditionalizing on a shared stream of evidence. That is, *inquiry* brings about agreement and decreases indeterminacy. Merging of opinions has been extended both to updating rules besides standard Bayesian conditionalization and certain large sets or priors. I bring these extensions together in

Chapter 4, showing that, under a few constraints, Jeffrey conditioning leads to merging of opinions for certain large sets.

The second section treats full beliefs in the setting of belief revision theory. Belief revision theories deal with full beliefs on their own terms, without presupposing a bridge or reduction to degrees of belief as is recently in fashion. Given the role consensus is supposed to play in the social aspects of inquiry and deliberation, it is important that we can always identify a consensus as the basis of joint inquiry and deliberation. But in Chapter 5, I show that if we think of an agent revising her beliefs (genuinely or for the sake of the argument) to reach a consensus, then on the received view of belief revision, AGM belief revision theory, certain simple and compelling consensus positions are *not* always available.

The account of consensus that Chapter 5 deals with, *shared agreement*, is a conservative account of consensus in full belief. In a sense, it is the analogue of the conciliatory position in the literature on peer disagreement for *sets* of belief. A shift to consensus as shared agreement suspends judgment on beliefs about which parties do not agree. Isaac Levi is among those who advocate it, and in Chapter 6, I present a constructive challenge to articulate motivations for that account in terms of Levi's own theory of justified belief change: if the standard (epistemic) goals that justify revision do not apply to shifting to consensus *as he claims*, how is shifting to consensus in full belief ever justified? I go on to argue that consensus not only fails to motivate an affirmative answer to the question of *whether* to revise, but revising to consensus violates principles of *how* to revise (on Levi's preferred account). In particular, postulates of *mild contraction* are generally violated in revising to consensus. I argue that the most promising response to the challenge is that, in general, a revision to consensus as shared agreement (when rational) should be thought of as a hypothetical revision, not a genuine one. Hypothetical or *for the sake of the argument* revision is not subject to the same constraints as genuine revision.

I turn to preference aggregation in the third section. Like other aggregation frameworks, social choice theory is famously beset by some very substantial limitative results. Results like Arrow's Impossibility Theorem and Sen's Liberal Paradox are sometimes thought to be devastating for the prospect of reasonable, democratic collective choice. I propose a framework for consensus in preference that allows for *indeterminacy* and show that we can skirt both Arrow's and Sen's results. In other words, the framework allows more desiderata for preference aggregation to be jointly satisfied than determinate aggregation methods consistently can. Furthermore, the proposed account offers a way of resolving not only some important conceptual tensions in democratic theory between social choice theory and deliberative democracy, but also some tensions in welfare economics between consequentialist and procedural constraints on aggregation. Finally, the account admits philosophical motivation in terms of the same conception of neutrality between rival points of view that underwrites my approach to probability pooling and belief aggregation.

CHAPTER 2

Probabilistic Opinion Pooling with Imprecise Probabilities

1. Introduction

The problem of opinion aggregation is "the problem of determining a sensible formula for representing the opinions of a group" (Genest et al., 1986). Representations of group opinion are important in a number of contexts, from scientific advisory panels (on climate change, for example), to joint efforts in scientific inquiry, to decision making in various kinds of groups. In a Bayesian setting, group consensus is particularly important from a theoretical standpoint. The received view is that probabilistic opinions are *subjective* (Ramsey, 1990b; Savage, 1954; de Finetti, 1964). Forms of intersubjective agreement have been sought to replace the surrendered notion of objectivity (Genest and Zidek, 1986; Nau, 2002). Probabilistic opinion pooling is one proposal for finding such consensus. It is widely assumed that, for a group of Bayesians, a representation of group opinion should take the form of a (single) probability distribution. The central position of this essay is that, in certain philosophically interesting and important cases, such an assumption is not always appropriate.

At the end of their review article on pooling, Dietrich and List mention other approaches that "redefine the aggregation problem itself" (2014, p. 20). According to them, one such approach is the aggregation of imprecise probabilities.[1] Of the few accounts of aggregating probabilities that

A version of this chapter is published as a paper, coauthored with Ignacio Ojea Quintana, under the same title in the *Journal of Philosophical Logic*.

[1]Here I use *IP* as a general term, abstracting from the important distinction Isaac Levi makes between what he calls *imprecise* and *indeterminate* probability, or what Walley calls the *Bayesian sensitivity analysis* and *direct* interpretations, respectively. Roughly speaking, according to the first interpretation, while an agent is normatively committed to or descriptively in a state of numerically precise judgments of credal probability, these precise judgments may not be precisely elicited or introspected. On the second interpretation, imprecision is a feature of the credal

deal with imprecision, many tend to focus on cases in which the individual opinions are already imprecise. And such accounts do not proceed by generalizing the pooling framework, axioms, etc. (Moral and Del Sagrado, 1998; Nau, 2002). A general account of probabilistic consensus should cover cases in which probabilities are imprecise at the level of the individual (a topic to which I return towards the end of the essay). However, our aim is to call into question the assumption that group opinion should be represented by a single probability distribution when precision holds at the level of the individuals. In this effort, I extend a line of argument that uses limitative results concerning aggregation—results demonstrating the impossibility of jointly satisfying a set of formal pooling criteria for precise aggregation methods—as a springboard into IP (Walley, 1982; Seidenfeld et al., 1989). That is, the limitations of precise pooling motivate IP in the sense that certain IP models *do* satisfy desiderata for "group" opinion that precise models do not.

After presenting the basic mathematical framework for probabilistic opinion pooling, I review some of the central limitative results (Section 2). One contribution of the present essay is generalizing the pooling framework, framing pooling with imprecise probabilities in the mathematical language common in research on probability aggregation with precise probabilities (Sections 4 and 5). The particular IP model that I primarily focus on in this dissertation, as a proof of concept, is presented in Section 4 (in a sense that will be made clear and precise, our case for considering IP models of pooling does *not* rise and fall with this particular format). Even in cases in which individual probabilities are precise, demanding that the output of an aggregation method be a single probability function is overly restrictive. As I show, representations of group opinion in terms of sets of probability functions have some very nice features. On the one hand, IP allows for a plausible philosophical account of rational consensus (Section 3). On the other hand, the construction I study satisfies a number of the central pooling axioms that are not jointly satisfied

state itself and is not attributable to imperfect elicitation or introspection. It is possible, of course, for a credal state to be imprecise in both senses, that is, an indeterminate credal state could be incompletely elicited.

by any of the standard, precise pooling recipes on pain of triviality (Sections 5 and 6). I close by considering some potential objections (Section 7).

2. Pooling

A general framework for aggregating the probabilistic opinions of a group to form a collective opinion is that of *pooling*. Formally, a pooling method for a group of n individuals is a function

$$F : \mathbb{P}^n \to \mathbb{P}$$

mapping profiles of probability functions for the n agents (or simply the n distributions under consideration), $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, to *single* probability functions intended to represent group opinion, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. The probabilities are assigned to events, which we represent as subsets of a sample space, Ω. We assume that Ω is countable. The *agenda*, or the set of events under consideration, is assumed to be an algebra $\mathscr{A}$ of events over Ω, that is, a set of subsets of Ω closed under complementation and finite unions (in the general case, closure under countable unions yields a σ-algebra).[2] A function $\boldsymbol{p} : \Omega \to [0, 1]$ is a *probability mass function* (pmf) *iff* $\sum_{\omega \in \Omega} \boldsymbol{p}(\omega) = 1$. Abusing notation, we can define a probability *measure*, $\boldsymbol{p}$, on general events for a given pmf by $\boldsymbol{p}(E) = \sum_{\omega \in E} \boldsymbol{p}(\omega)$. Pooling can be formulated in terms of pmfs, and we will appeal to pmfs in discussing geometric pooling functions and the external Bayesianity constraint below.

Various interpretations of pooling are proposed in the literature. Wagner, for example, offers the following (2009, pp. 336-337):

(1) A rough summary of the current probabilities of the n individuals;

[2]For completeness, I include the probability axioms. A *probability function* is a mapping $\boldsymbol{p} : \mathscr{A} \to \mathbb{R}$ that satisfies the following conditions:

(i) $\boldsymbol{p}(A) \geqslant 0$ for any $A \in \mathscr{A}$;
(ii) $\boldsymbol{p}(\Omega) = 1$;
(iii) $\boldsymbol{p}(A \cup B) = \boldsymbol{p}(A) + \boldsymbol{p}(B)$ for any $A, B \in \mathscr{A}$ such that $A \cap B = \varnothing$.

If, in addition, $\mathscr{A}$ is a σ-algebra and $\boldsymbol{p}$ satisfies the following condition, $\boldsymbol{p}$ is called *countably additive*:

(iv) If $\{A_n\}_{n=1}^{\infty} \subseteq \mathscr{A}$ is a collection of pairwise disjoint events, then $\boldsymbol{p}(\bigcup_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} \boldsymbol{p}(A_n)$.

In this essay, I assume countable additivity for convenience, not because I take it to be rationally mandatory.

(2) a "compromise" adopted by the individuals for the purpose of group decision making;

(3) a rational consensus to which the individuals revise their probabilities after discussion;

(4) the opinion a decision maker external to the group adopts upon being informed of the n expert opinions in the group;

(5) the opinion an individual in the group adopts upon being informed of the $n-1$ opinions of his "epistemic peers" in the group.

These five interpretations do not exhaust the possibilities. Our target interpretation is rational consensus, adopted either *for the sake of the argument* (a compromise) in order to perform some task in group inference or decision making (2) or genuinely by individual group members (3, 5). However, the account I consider could also be used by a decision maker external to the group.

2.1. Criteria for Pooling Functions. What properties should a pooling function have? Let's review some of the most popular properties discussed in this connection. It is important to consider, for each property, the extent to which it is normatively compelling for a particular interpretation and use of pooling functions. Surveys of the material presented here include Simon French's (1985), Genest and Zidek's (1986), and Dietrich and List's (2014).

McConway (1981) and Lerher and Wagner (1981) introduce a convenient property of pooling functions called the *strong setwise function property* and *strong label neutrality* by the respective authors. The property has it that the individual probabilities for an event—and not the entire distributions of each individual—are all that is required to determine the collective probability of that event.

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

What case can be made for the strong setwise function property (SSFP) as a pooling *norm*? SSFP can be seen as a probabilistic analogue of the independence of irrelevant alternatives constraint in the social choice literature. Consider two profiles $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ and $(\boldsymbol{p}'_1, ..., \boldsymbol{p}'_n)$. Suppose that, for some event A, $\boldsymbol{p}_i(A) = \boldsymbol{p}'_i(A)$ for $i = 1, ..., n$, but the two profiles differ on some other event (so for pooling probabilities for A, "irrelevant" parts of the probability functions differ). It can happen that $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) \neq F(\boldsymbol{p}'_1, ..., \boldsymbol{p}'_n)(A)$ despite the fact that $(\boldsymbol{p}_1(A), ...\boldsymbol{p}_n(A)) = (\boldsymbol{p}'_1(A), ..., \boldsymbol{p}'_n(A))$. That is, the "consensus" probabilities for A differ for the two profiles despite no change in individual opinions concerning A. So, the pooled probability for A is not a function merely of the individual probabilities for A. For such an F, no function G exists because such a function would have to map one profile of values in $[0, 1]^n$ to two distinct outputs in $[0, 1]$. Admittedly, such a case for the normative status of SSFP is incomplete.

Many of the axioms proposed in the literature on pooling require that some property of the individual probability functions be preserved under pooling. When the algebra contains at least three events, one such preservation property follows immediately from SSFP, as McConway observes (1981, Theorem 3.2).

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

As Genest and Zidek remark, the zero preservation property (ZPP) is one in a class of constraints requiring that the pool preserves initial shared agreements. The normative status of this sort of preservation axiom has been called into question in the literature (Genest and Wagner, 1987). Of course, ZPP is forced upon those endorsing SSFP. For conceptions of *consensus* on which common ground is sought, that is, a non-question begging position of agreement, ZPP is more compelling. I return to *consensus as shared agreement* or *common ground* below in Section 3.

McConway's Theorem 3.2 shows more. Taken together, the *marginalization property* (MP) and the zero preservation property (ZPP) are equivalent to SSFP. McConway's formal setup differs somewhat from the one presented here. He is concerned with classes of pooling functions that take into account all $\sigma-$algebras on Ω. We, however, are considering pooling functions for a fixed algebra (which seems to be the more common approach). The formal properties of concern to McConway must be modified accordingly. A pooling function satisfies MP if marginalization and pooling commute. We adopt the modification of MP proposed by Genest and Zidek (1986, p. 118). Let $\mathscr{A}'$ be a subalgebra of $\mathscr{A}$.[3] Suppose that $\boldsymbol{p}$ is a distribution over $(\Omega, \mathscr{A})$. The *marginal* distribution $\boldsymbol{p}\restriction_{\mathscr{A}'}$ given by $\boldsymbol{p}$ over $(\Omega, \mathscr{A}')$ is the restriction of $\boldsymbol{p}$ to $\mathscr{A}'$ such that $\boldsymbol{p}(A) = \boldsymbol{p}\restriction_{\mathscr{A}'}(A)$ for all $A \in \mathscr{A}'$. $[\boldsymbol{p}\restriction_{\mathscr{A}'}]$ is a Carathéodory extension of $\boldsymbol{p}\restriction_{\mathscr{A}'}$ to $\mathscr{A}$.

Marginalization Property. Let $\mathscr{A}'$ be a sub-σ-algebra of $\mathscr{A}$. For any $A \in \mathscr{A}'$, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = F([\boldsymbol{p}_1\restriction_{\mathscr{A}'}], ..., [\boldsymbol{p}_n\restriction_{\mathscr{A}'}])(A)$.

Below, I will state an analogue of another of McConway's results. That result says that MP is equivalent to the *weak setwise function property* (WSFP) (1981, Theorem 3.1). Instead of $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$ depending just on the $\boldsymbol{p}_i(A), i = 1, ...n$, those pooling functions merely satisfying WSFP depend on both $\boldsymbol{p}_i(A)$ *and* the event, A. The difference is that a profile in $[0, 1]^n$ may be mapped to more than one output, so long as the associated event differs.

Weak Setwise Function Property. There exists a function $G : \mathscr{A} \times [0, 1]^n \to [0, 1]$ such that, for any event $A \in \mathscr{A}$, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = G(A, \boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A))$ for each profile in the domain of F.

Probabilistic independence is another natural candidate property for preservation under pooling. In the precise setting, there are a number of equivalent formulations of probabilistic independence. For example, two events, A and B, are said to be *stochastically independent* according to $\boldsymbol{p}$ if $\boldsymbol{p}(A \cap B) = \boldsymbol{p}(A)\boldsymbol{p}(B)$. Dividing both sides by $\boldsymbol{p}(B)$, provided $\boldsymbol{p}(B) > 0$, yields $\frac{\boldsymbol{p}(A \cap B)}{\boldsymbol{p}(B)} = \boldsymbol{p}(A)$ when A and B are independent. But the lefthand side of the equation is a standard definition of the probability of A *conditional* on B: $\boldsymbol{p}(A|B) = \frac{\boldsymbol{p}(A \cap B)}{\boldsymbol{p}(B)}$, when $\boldsymbol{p}(B) > 0$. This observation allows us to state another standard formulation of probabilistic independence. A and B are independent according to $\boldsymbol{p}$ if $\boldsymbol{p}(A|B) = \boldsymbol{p}(A)$. The *conditionalization* of $\boldsymbol{p}$ with respect to an event B, $\boldsymbol{p}^B$, is

[3] $\mathscr{A}'$ is a boolean subalgebra of $\mathscr{A}$ if $\mathscr{A}' \subseteq \mathscr{A}$ and $\mathscr{A}'$, with the distinguished elements and operations of $\mathscr{A}$, is a boolean algebra. That is, the operations must be the restrictions of the operations of the whole algebra; being a subset that is a boolean algebra is not sufficient for being a subalgebra of $\mathscr{A}$ (Halmos, 1963).

given by setting $\boldsymbol{p}^B(A) = \boldsymbol{p}(A|B)$ for all A. I will return to stochastic independence below, but it will be convenient for us to adopt the definition in terms of conditional probabilities.

Probabilistic Independence Preservation. If $\boldsymbol{p}_i(A|B) = \boldsymbol{p}_i(A)$ for $i = 1, ..., n$, then $F^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$.

This axiom says that two events that are probabilistically independent according to every individual probability function are independent according to the pool.

Another preservation axiom is *unanimity preservation*, which requires that, if all of the functions being pooled are identical, then the output of the pooling function is that probability function. So if all the individual opinions are the same, the group opinion is identical to that common distribution.

Unanimity Preservation. For every opinion profile $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$, if all $\boldsymbol{p}_i$ are identical, then $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \boldsymbol{p}_i$.

Other sorts of pooling axioms, like MP above, demand that some operation or other commutes with pooling. A very interesting example of such an operation is a type of Bayesian updating. Standard Bayesian conditionalization goes *via* Bayes' theorem:

$$\boldsymbol{p}^B(A) = \boldsymbol{p}(A|B) = \frac{\boldsymbol{p}(A)\boldsymbol{p}(B|A)}{\boldsymbol{p}(B)}, \text{ when } \boldsymbol{p}(B) > 0.$$

By the law of total probability, the denominator, $\boldsymbol{p}(B)$, can be rewritten. Where $\{C_j : j = 1, 2, ...\}$ is a partition of Ω, $\boldsymbol{p}(B) = \sum_j \boldsymbol{p}(B|C_j)\boldsymbol{p}(C_j)$.

External Bayesianity is a mild generalization of commutativity with Bayesian conditionalization. The requirement is that updating the individual probabilities on a common *likelihood function* (as opposed to updating on an event) and then pooling is the same as pooling and then updating the pool on that likelihood function. The likelihood function, $\lambda : \Omega \to [0, \infty)$, is defined on elements of the sample space. In conditionalizing, $\lambda(\cdot)$ serves the same role as the conditional probability $\boldsymbol{p}(B|\cdot)$ in Bayes' theorem above, expressing the degree to which some fixed evidence B is expected on various events. Put roughly, updating on λ results from substituting the likelihood function in for the conditional probabilities on the right hand side of Bayes' theorem. For every $\omega \in \Omega$,

$$\boldsymbol{p}^{\lambda}(\omega) = \frac{\boldsymbol{p}(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')}, \text{when } \sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega') > 0$$

If $\boldsymbol{p}$ is a probability measure, it must be defined on an algebra including the elements of Ω. Otherwise, take $\boldsymbol{p}$ to be a pmf and obtain a probabiltiy measure on a given algebra by summing over the elements of Ω in each event to obtain the probability of events in the algebra. Comparing the above formula with the version of Bayes' theorem in which the denominator is expanded by the law of total probability makes the relation between $\lambda(\omega)$ and $\boldsymbol{p}(B|\cdot)$ apparent. While not itself a probability distribution, $\lambda(\omega)$ is proportional to $\boldsymbol{p}(B|\omega)$, for fixed data B. And though not a function of general events in $\mathscr{A}$, the likelihood of an event A can be obtained by summing the likelihoods of all $\omega \in A$. Updating on a likelihood function reduces to standard conditionalization on some event, B, when

$$\lambda(\omega) = \begin{cases} 1, & \text{if } \omega \in B \\ 0, & \text{otherwise.} \end{cases}$$

(We verify the claim with routine subsititions in the footnote.[4])

Crucially, the likelihood function is assumed to be *common* in the external Bayesianity axiom. So while disagreement concerning the prior is permitted by pooling functions satisfying external Bayesianity, the commutativity of pooling and updating is guaranteed only when there is agreement on the likelihood funciton.

[4]For any $A \in \mathscr{A}$, $\boldsymbol{p}^E(A) = \frac{\boldsymbol{p}(A \cap E)}{\boldsymbol{p}(E)} = \frac{\sum_{\omega \in A \cap E} \boldsymbol{p}(\omega)}{\sum_{\omega \in E} \boldsymbol{p}(\omega)}$. By the definition of a probability measure, $\boldsymbol{p}(A) = \sum_{\omega \in A} \boldsymbol{p}(\omega)$, so $\sum_{\omega \in A} \boldsymbol{p}^{\lambda}(\omega) = \frac{\sum_{\omega \in A} \boldsymbol{p}(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')}$ gives us $\boldsymbol{p}^{\lambda}(A)$. We show that these two fractions are equal by showing the equality of both the numerators and denominators. Since, for all $\omega \in A$, $\boldsymbol{p}(\omega)\lambda(\omega) = \boldsymbol{p}(\omega)$ if $\omega \in E$ and 0 otherwise, $\sum_{\omega \in A} \boldsymbol{p}(\omega)\lambda(\omega) = \sum_{\omega \in A \cap E} \boldsymbol{p}(\omega) = \boldsymbol{p}(A \cap E)$. Hence, the numerators are equal. And since, for all $\omega' \in \Omega$, $\boldsymbol{p}(\omega')\lambda(\omega') = \boldsymbol{p}(\omega')$ if $\omega' \in E$ and 0 otherwise, we have $\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega') = \sum_{\omega' \in E} \boldsymbol{p}(\omega') = \boldsymbol{p}(E)$. Hence, the denominators are equal, too. So, $\boldsymbol{p}^E = \boldsymbol{p}^{\lambda}$.

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

A similar axiom requires that a single individual conditionalizing on λ before pooling is the same as conditionalizing the pool on λ (Dietrich and List, 2014).

Individualwise Bayesianity. For every profile $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ in the domain of F and every individual k such that $(\boldsymbol{p}_1, ..., \boldsymbol{p}_k^\lambda, ..., \boldsymbol{p}_n)$ remains in the domain, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_k^\lambda, ..., \boldsymbol{p}_n) = F^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$

I also have more to say about individualwise Bayesianity below.

2.2. Types of Pooling Functions. Various concrete pooling functions have been studied in the literature. These functions fare differently on the criteria reviewed just above. Of the commonly discussed pooling operators, linear pooling functions may be the most common and obvious proposal (Stone, 1961; McConway, 1981; Lehrer and Wagner, 1981).

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

$w_1, ..., w_n$ are fixed non-negative weights summing to 1 that are associated with the n individuals. Linear pooling, then, takes a weighted average of the individual probabilities. Equal weights for the n probability functions specifies one linear pooling function; a *dictatorship* specifies another linear pooling function. In the latter case, all of the weight is accorded to a single individual ($w_i = 1$ for some i) with the result that the pooled probability for any event A is that individual's probability for A: $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \boldsymbol{p}_i(A)$. Interestingly, weights $w_i = \frac{1}{n}$ were used in a U.S. Nuclear Regulatory Commission study of the frequency of nuclear reactor accidents (Ouchi, 2004, p. 5).

Another proposal is to take a weighted *geometric* instead of a weighted arithmetic average of the n probability functions (Madansky, 1964; Bacharach, 1972; Genest et al., 1986).[5]

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

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

Unlike linear pools, geometric pools specify the collective probabilities of elements of Ω instead of events in general. But as with the likelihood functions above, the probability of any event A is determined by summing the probabilities of $\omega \in A$. Because of the way in which multiplication figures into the geometric pooling recipe, there are profiles for which $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(\omega) = 0$ for all $\omega \in \Omega$, in violation of the probability axioms. If for each $\omega \in \Omega$ there is a $\boldsymbol{p}_i \in (\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ such that $\boldsymbol{p}_i(\omega) = 0$ we have such a violation. To avoid this worry, the domain of geometric pooling operators is restricted to profiles of *regular* pmfs, i.e., those $\boldsymbol{p}$ such that $\boldsymbol{p}(\omega) > 0$ for all $\omega \in \Omega$. We denote the set of regular pmfs $\mathbb{P}'$ making the relevant domain $\mathbb{P}'^n$.[6]

A third, more recent proposal from Dietrich (2010) is given by the following formula.

Multiplicative Opinion Pools. $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(\omega) = c\prod_{i=0}^{n} \boldsymbol{p}_i$, where $\boldsymbol{p}_0$ is a fixed "calibrating" probability function, and $c = \frac{1}{\sum_{\omega' \in \Omega}[\boldsymbol{p}_0(\omega)] \cdot [\boldsymbol{p}_1(\omega')] \cdots [\boldsymbol{p}_n(\omega')]}$ is a normalization factor.

As with geometric pooling functions, the domain of multiplicative pooling functions will be restricted to $\mathbb{P}'^n$. Comments on the interpretation and choice of $\boldsymbol{p}_0$ can be found in (Dietrich and List, 2014, pp. 17-19)

Various results, both characterization and limitative, for the different pooling operations and axioms have been obtained. For example, SSFP characterizes linear pooling.

THEOREM 1. *(McConway, 1981, Theorem 3.3; Lehrer and Wagner, 1981, 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.*

[6] Rather than assuming regularity or that the algebra contains the elements of Ω, we could make the weaker restriction to the domain of profiles of pmfs such that there is some $\omega \in \Omega$ for which $\boldsymbol{p}_i(\omega) > 0$ for all $i = 1, ..., n$. And pmfs allow us to obtain measures defined on general algebras on Ω.

Theorem 2. *(McConway, 1981, 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. *(Genest, 1984, p. 1104) The geometric pooling functions are externally Bayesian and preserve unanimity.*

Other sorts of pooling functions, such as a certain generalization of geometric pooling, satisfy the conditions of Theorem 3. A characterization of externally Bayesian pooling functions is given in (Genest et al., 1986). Dietrich and List provide a characterization of multiplicative pooling.

Theorem 4. *(Dietrich and List, 2014, Theorem 3) The multiplicative pooling functions are the only individualwise Bayesian pooling functions (with domain $\mathbb{P}'^n$).*

There are many Arrovian limitative theorems in the pooling literature. As Robert Nau notes, none of the pooling methods discussed satisfy even unanimity, external Bayesianity, and the marginalization property (2002, p. 266). (As I show below, our proposal in this essay *does* satisfy those properties.) One result that we have occasion to appeal to below follows from results due to Lehrer and Wagner (1983, Theorems 1 and 2) in conjunction with Theorem 1 above:

Theorem 5. *(Cf. Lehrer and Wagner, 1983) Given that the algebra contains at least three pairwise disjoint events, the only pooling functions that preserve probabilistic independence and satisfy SSFP are dictatorial.*

It follows that non-dictatorial *linear pools* do not preserve probabilistic independence. In general, non-dictatorial pooling methods struggle with independence preservation (Genest and Wagner, 1987). (Here, too, I claim to do better.)

3. Motivations for IP

In general terms, imprecise probabilities (IP) models do not require representing an agent's or group's judgments of subjective probability as numerically precise. Instead, such judgments could be represented by *sets* of probability functions (Kyburg and Pittarelli, 1996), for example, or by *intervals* (Kyburg, 1998).

There are a number of motivations for working with IP models. These include the potential to resolve some of the "paradoxes of decision" (Ellsberg, 1963; Levi, 1986b), allowing for more flexible and less arbitrary models of uncertainty (Gärdenfors and Sahlin, 1982; Walley, 1991), allowing for incomplete preferences (and hence judgments of incomparability) in the subjective expected utility setting (Levi, 1986a; Seidenfeld, 1993; Kaplan, 1996), and increased descriptive realism (Arló-Costa and Helzner, 2010). An overview of these and other motivations for IP can be found in (Bradley, 2014).

Most important for present purposes, IP allows for—what I consider—a very interesting and philosophically well-motivated account of consensus (Levi, 1985; Seidenfeld et al., 1989). Our goal in this section is to present this account of consensus for explicit consideration in the context of pooling. It may help to first consider the case of *full* or *plain* belief. At the outset of inquiry, inquirers may seek consensus as *shared agreement* in their beliefs. This could be achieved by retaining whatever beliefs are common to all parties and suspending judgment on those that are controversial thereby avoiding question-begging. Importantly, the consensus is generally a *weaker* state of belief. Since inquiry initiating from the consensus view proceeds without begging questions against parties to the consensus, various hypotheses of concern can receive a fair hearing. Such a consensus constitutes a neutral or non-controversial starting point for subsequent inquiry.

The idea that parties to a joint effort in inquiry or decision making should restrict themselves to their shared agreements—as a compromise or as genuine consensus—can be extended to judgments of probability. An analogous sense of suspending judgment concerning what is controversial is available in the IP setting. To suspend judgment among some number of probability distributions is to not rule them out for the purposes of inference and decision making. Put another way, to suspend judgment among some number of distributions is to regard each as *permissible* to use in inference and decision making. If the parties seeking consensus all agree that $\boldsymbol{p}$ is *not* permissible, then the consensus position reflects that agreement and rules it out (this will have to be finessed when we come to the question of convexity below). A *set* of probability functions represents the shared agreements among the group concerning which probability functions *are not* permissible to use in inference and decision making. For example, it is consensus that the probability of some event is not below the minimum of individual assignments.

Many authors refer to the output of a pooling function as a *consensus* (Lehrer and Wagner, 1981; McConway, 1981; Genest and Zidek, 1986). In what way is a precise pool a consensus? Isaac Levi draws a distinction between consensus as the *outcome* of inquiry and consensus at the *outset* of inquiry (1985). At the outset of inquiry, agents may seek common ground upon which to pursue joint inquiry. This is consensus as shared agreement, discussed just above. Disagreement among the parties to the consensus may then be resolved (in the best case) through joint efforts in inquiry—consensus as the outcome of inquiry. Given the restriction that consensus must be representable by a unique probability function, outside of the special case when all individuals are in total agreement, finding consensus as shared agreement that suspends judgment on unshared probabilistic views is a foreclosed possibility.

The individuals could assume some common, precise probability distribution, but Levi argues this is not consensus as common ground:

> there can be no analogue in contexts of probability judgment of the two senses of consensus I identify. If two or more agents differ in probability judgment, they can all switch either to the distribution adopted by one of them or to some other distribution which is, in a sense, a potential resolution of the conflict between their differing distributions. There is only one kind of consensus to be recognized—namely the resolution of conflict reached through revolution, conversion, voting, bargaining or some other psychological or social process. (1985, pp. 5-6)

Wagner likewise distinguishes between a *compromise* adopted to perform an exercise in group decision making and a *consensus* to which the individuals revise their own beliefs (Section 2). Levi's point in the quotation above is that a precise pool is neither a consensus as shared agreement since, in general, it is not restricted to just the shared probabilistic views; nor is it justified on the basis of inquiry, understood as designing and performing experiments, obtaining evidence, etc. A precise pool might represent the sort of political consensus that a vote does in the case of preferences, or that the output of a judgment aggregation function does in the case of qualitative belief. That is, consensus as bargaining or compromise. Of course, at least one sort of compromise *is* a consensus adopted *for the sake of the argument* rather than genuinely. That is, parties to the compromise could assume the consensus position as Levi identifies it—namely, a convex IP set—for the sake of the argument, or for carrying out some group deliberation or inquiry so long as the consensus position is strong enough for the group's purposes. It must be admitted that there are other compromise positions, including precise pools, that the group might assume.

But Levi's view distinguishes between political and rational consensus. Returning again to the case of full belief, Levi requires that revisions be decomposable into a sequence of contractions and expansions. An inquirer's set of full beliefs constitute her "standard for serious possibility" in the following sense: if A is among her full beliefs, $\neg A$ is not a serious possibility. To change her mind, the inquirer must first suspend judgment on A by contracting A if she cares to avoid error (where error is judged by her own lights). From the contraction, both A and $\neg A$ are serious possibilities. A *direct* revision to include $\neg A$ involves deliberately importing error from the point of view that rules $\neg A$ out as a serious possibility.

Unlike full beliefs, however, judgments of subjective probability do not bear truth values. So how might one suspend judgment among candidate distributions before changing points of view? As discussed just above, subjective probabilities are used in determining expectations for available acts. Levi's proposal is that to suspend judgment among some number of distributions is to not rule them out for use in the functions that they perform in inquiry and deliberation. Coming to regard a probability distribution as *permissible* is the analogue of opening one's mind to the (serious) *possibility* of $\neg A$ in the case of full belief. Just as assuming a weaker position in full belief avoids begging questions, retreating to a superset of distributions avoids prejudging the issue of determining which distributions are permissible for use in inquiry and deliberation. Moving from a set of probability functions (including a singleton) to a superset is the probabilistic analogue of contraction. Employing sets of probability functions avoids demanding direct revisions to probabilistic judgments the agent regards as *impermissible* from the standpoint of her current probabilistic judgments without first "contracting" to a neutral position that suspends judgment among the relevant probabilistic views (Levi, 1974). As he emphasizes, both reaching common ground at the outset of inquiry and subsequent reasoned changes in probabilistic views are available to groups in the IP setting.

Why has most work on probabilistic aggregation restricted itself to consideration of representations in terms of a single probability function? One reason is that such representation is the standard for individuals and, since we are treating groups as agents in a sense, that representation should extend to groups as well. Genest and Zidek write, "it would be natural to express the consensus judgment in the same form as the originals" (1986, p. 115). But I am not moved by this convention (or in Walley's terminology, by this "Bayesian dogma of precision") for some of the very reasons as discussed just above. I urge, in what follows, that theorizing concerning IP should be extended to accounts of probabilistic opinion pooling and *vice versa*. Even if the motivations for IP in general, including at the level of individuals, are found less than convincing, one might think that the case for IP at the level of group opinion is more persuasive, say as an account of consensus. I take the motivations above and the propositions that follow as recommending further consideration of IP in the context of pooling and consensus.

4. IP Pooling Formats

I want to make a case that the out-of-the-gate restriction of the codomain of F to $\mathbb{P}$ is unwarranted (just as many have argued that the standard Bayesian assumption that rational individuals are committed to determinate probabilistic judgments is unwarranted). Our strategy is to point to a sensible account that abandons that restriction. Here we assume a representation in terms of a set of probability functions. We make use of *set-valued* functions or *correspondences*. Where F refers to a pooling function that outputs a single probability function, we will use $\mathcal{F}$ to refer to pooling correspondences outputting sets of probability functions:

$$\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$$

4.1. Convex or Not? Much of the work with IP assumes that IP sets of probabilities are *convex* (Smith, 1961; Levi, 1974; Girón and Ríos, 1980; Gilboa and Schmeidler, 1989; Walley, 1991; Moral and Del Sagrado, 1998). A set of probability functions, $\boldsymbol{P}$, is convex if, for any two functions in the set, the set includes every convex combination of those functions.

Convexity. If $\boldsymbol{p}_1, \boldsymbol{p}_2 \in \boldsymbol{P}$, then $\alpha\boldsymbol{p}_1 + (1-\alpha)\boldsymbol{p}_2 \in \boldsymbol{P}$ for $\alpha \in [0,1]$.

Besides some handy computational properties of convex sets of probability functions (Girón and Ríos, 1980), convexity can be motivated philosophically. A set of probability functions represents the shared agreements among the group members concerning which distributions are ruled out for use in deliberation and inquiry. But is it not common ground that the convex combinations of individual probability functions are ruled out? The idea is that convexity recommends a *weaker* attitude in suspending judgment among some number of probability distributions; fewer distributions are ruled out. Convexity requires keeping an open mind concerning potential compromises or resolutions of conflict (the convex combinations) between various probabilistic views. Levi argues that convex combinations have "all the earmarks of potential resolutions of the conflict; and, given the assumption that one should not preclude potential resolutions when suspending judgment between rival systems [...] all weighted averages of the two functions are thus taken into account" (1980, p. 192).

The normative status of convexity is the subject of outstanding controversy. Seidenfeld, Schervish, and Kadane make a case against convexity in the context of group decision making (1989). They observe that if two Bayesian agents differ in both probability and utility, any compromise position in probability besides *trivial* convex combinations entails a violation of a Pareto constraint on preference. Levi responds in his (1990), arguing against the Pareto condition. Kyburg and Pittarelli lodge some complaints about the property in "Some Problems for Convex Bayesians" (1992). In "Set-Based Bayesianism," they explore relaxing convexity to allow for IP sets in general. Seidenfeld et al.'s theory of coherent choice does not require convexity (2010). They offer a variation of one of Kyburg and Pittarelli's criticisms of convexity, registering a counterexample that exploits the failure of convex combinations to preserve probabilistic independence (but see (Levi, 2009, pp. 373-375) for a response).

Depending on the decision theory, distinctions between pooling formats may or may not be of importance. For example, there are decision rules that cannot distinguish between certain convex and non-convex sets of probabilities (Gilboa and Schmeidler, 1989; Walley, 1991). Such distinctions are meaningful according to other decision rules (Levi, 1980). And there are decision rules that distinguish between any two sets of probabilities (Seidenfeld et al., 2010). The important point here is that disputes over the format of pooling functions are idle if such distinctions are not decision-theoretically meaningful. On decision theories that cannot distinguish a convex set of probabilities from its extreme points, for example, there is nothing at stake in arguments over whether an IP opinion pool is convex or not.

4.2. Convex Pooling Functions. As a proof of concept, I will investigate aggregation functions that output sets of probability functions. The aggregate is formed by taking the *convex hull* of the n probability distributions:

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv\{\boldsymbol{p}_i : i = 1, ..., n\}$$

The convex hull of a set of points is the smallest convex set containing those points. We write $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$ as shorthand for the set of probability assignments to A:

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$$

I work with convexity, not because I presume to know of decisive arguments in its favor, but because convexity is a broadly customary assumption, I do not yet feel compelled to dismiss it, and it allows me to make a proof of concept argument for IP pooling. In effect, assuming convexity amounts to making it slightly harder to show some of the propositions below, though the propositions also hold for IP aggregation methods that relax convexity. I return to the issue of convexity below to make good on my earlier promise to clarify how our case for IP in the context of opinion aggregation does not depend entirely on *convex* IP pools (Section 7.3, Proposition 7).

5. Extending Pooling Axioms to the IP Setting

Convex IP pooling functions satisfy the extensions of those axioms to the IP setting. For the SSFP, we replace G with a set-valued function or correspondence: $\mathcal{G} : [0,1]^n \rightarrow \mathscr{P}([0,1])$. $\mathcal{G}$ is a map from n numerical values in $[0,1]$ to a set of probability values, $\mathcal{G}(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot))$. The constraint becomes that there exists a function $\mathcal{G}$ such that, for any event A, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{G}(\boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A))$. WSFP, then, requires a function $\mathcal{G} : \mathscr{A} \times [0,1]^n \rightarrow \mathscr{P}([0,1])$. For unanimity preservation, we do not distinguish between $\boldsymbol{p}$ and $\{\boldsymbol{p}\}$. ZPP is generalized analogously. If $\boldsymbol{p}_i(A) = 0$ for all $i = 1, ..., n$, then $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \{0\}$. The MP has a straightforward extension to sets of probability functions: $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}([\boldsymbol{p}_1 \restriction_{\mathscr{A}'}], ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}])(A)$, for any $A \in \mathscr{A}'$ There are many ways to generalize constraints. While I offer conservative and natural modifications of the axioms in order to extend them to the imprecise setting, the crucial question is whether I have modified what is *compelling* about the axioms. For example, is representation in terms of a unique probability function crucial to what makes commutativity of conditionalization and pooling compelling, or what is appealing about preserving shared judgments of independence? In each case, I submit, the attractiveness of the axiom does not hinge on whether the output of the aggregation function is a single probability function or a set of them.

First, we note that an analogue of McConway's result holds for IP pooling functions in general.

PROPOSITION 1. *Let $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$ be an IP pooling function (not necessarily convex). $\mathcal{F}$ satisfies WSFP iff $\mathcal{F}$ satisfies MP.*

Before stating the next proposition, we record a fact about convex sets of probabilities (simple and familiar to those with a background in geometry) that we will make use of in the proof.[7] Proofs for the lemmas and propositions are recorded in the appendix.

LEMMA 1. *Let $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv\{\boldsymbol{p}_i : i = 1, ..., n\}$ for any profile $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ in the domain of $\mathcal{F}$. Any $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ can be expressed as a convex combination of the n probability functions, i.e. $\boldsymbol{p} = \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i$, where $\alpha_i \geqslant 0$ for $i = 1, ..., n$ and $\sum_{i=1}^{n} \alpha_i = 1$.*

PROPOSITION 2. *Convex IP pooling functions satisfy SWFP, WSFP, MP, ZPP, and unanimity preservation.*

As indicated in the proof, SSFP entails both WSFP and ZPP.

While linear pooling functions are not externally Bayesian, convex IP pooling functions satisfy the extension of external Bayesianity to the IP setting. The *convex* or *prior-by-prior* conditionalization of a convex set of probability functions, $\mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, results from conditionalizing each

[7] I include a proof of the observation because we appeal to it several times in the other proofs, because it is a simple special case (but all we need) of a more general result concerning convexity (Rockafellar, 1970), and because some readers may not have a conceptual handle on the property.

member of the set. Updating a convex set of probability functions on a common likelihood function is defined analogously:

$$\mathcal{F}^{\lambda}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \left\{ \boldsymbol{p}^{\lambda} : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n), \sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega') > 0, \text{ and } \boldsymbol{p}^{\lambda}(\cdot) = \frac{\boldsymbol{p}(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')} \right\}$$

To show that convex IP pooling functions are externally Bayesian, we first state another observation.

LEMMA 2. *(Cf. Levi, 1978; Girón and Ríos, 1980) Convexity is preserved under updating on a likelihood function, i.e.,* $\mathcal{F}^{\lambda}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ *is convex.*

PROPOSITION 3. *Convex IP pooling functions are externally Bayesian.*

Dietrich and List argue that while geometric pooling is justified on epistemic grounds when individual opinions are based on the same information, multiplicative pooling is justified in cases of *asymmetric* information, when individual opinions are in part based on private information. Their case is built around the individualwise Bayesianity axiom and the fact that multiplicative pooling satisfies it (Theorem 4).

PROPOSITION 4. *Convex IP pooling functions are **not** individualwise Bayesian.*

I regard Proposition 4, however, as stating a feature and not a bug of convex IP aggregation. At least insofar as the idea is to reach a consensus, it is not desirable for features of one individual's probability distribution to be unilaterally imposed on the group. In the case of full belief, the initial consensus does not adopt just any belief of any member. Better, in our view, for group opinion to change through efforts in intelligently conducted inquiry from initial common ground (in inquiry, a group may designate a subgroup as an information source on a given topic, but this process requires a richer representation). Dietrich and List motivate individualwise Bayesianity by pointing out that if the constraint is not satisfied, then it makes a difference if an individual first learns some information and opinions are then pooled, or if the opinions are pooled and then the information is acquired by the group as a whole. But for consensus, this is as it should be. If the opinions of the group members do not reflect some piece of information, that information is not common ground. The consensus among group members depends on the probabilistic opinions of the members.

None of this is to say, of course, that features of individual probability distributions are irrelevant to group consensus. On the convex IP view, the kernel of truth in individualwise Bayesianity can be formulated by the inequalities below, stated here for standard conditionalization.

$$\begin{aligned} & min\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\} \\ \leqslant\ & min\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_i^B, ..., \boldsymbol{p}_n)\} \\ \leqslant\ & min\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\} \end{aligned}$$

And similarly,

$$\begin{aligned} & max\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\} \\ \leqslant\ & max\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_i^B, ..., \boldsymbol{p}_n)\} \\ \leqslant\ & max\{\boldsymbol{p}(B) : \boldsymbol{p} \in \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}. \end{aligned}$$

The consensus probabilities for B shift up (at least not down, more precisely) if one individual conditionalizes on B, and shift more if the consensus itself conditionalizes on B. But these inequalities simply reflect facts about what the common ground is and do not reflect group "learning" from one individual's probability function.

Convex IP pooling also inherits some of the challenges facing linear pooling. SSFP conflicts with probabilistic independence preservation. As Theorem 5 states, the only pooling functions that preserve probabilistic independence and satisfy SSFP are dictatorial. The loss of probabilistic independence presents both epistemic challenges as well as decision theoretic ones (Kyburg and Pittarelli, 1992; Seidenfeld et al., 2010).

In the case of convex IP pooling, however, there is more leeway to address the challenges. Several generalizations of the concept of independence for IP have been proposed and studied (de Campos and Moral, 1995; Cozman, 1998). We consider Levi's notion of *confirmational irrelevance*.

Confirmational Irrelevance Preservation. If $\boldsymbol{p}_i(A|B) = \boldsymbol{p}_i(A)$ for $i = 1, ..., n$, then $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$.[8]

Irrelevance preservation is a generalization of probabilistic independence preservation. It is clear that when $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is a single probability function, irrelevance preservation reduces to independence preservation. According to some decision theories for IP, it is the whole set $\boldsymbol{P}$ that is relevant for inquiry and decision making (Levi, 1980; Seidenfeld et al., 2010). Irrelevance is a sensible generalization of independence because it allows us to identify when some information will not make a difference to certain inquiries or deliberations, namely, those inquiries and deliberations concerning events to which the information is irrelevant.

It also does not take much work to show that confirmational irrelevance preservation is satisfied by any IP pooling function (not necessarily convex) that satisfies *stochastic independence preservation.*

Stochastic Independence Preservation If $\boldsymbol{p}_i(A \cap B) = \boldsymbol{p}_i(A)\boldsymbol{p}_i(B)$, for $i = 1, ..., n$, then, for all $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, $\boldsymbol{p}(A \cap B) = \boldsymbol{p}(A)\boldsymbol{p}(B)$.

It turns out that confirmational irrelevance *is*, but stochastic independence is *not*, preserved by convex IP pooling functions. Suppose that A and B are probabilistically independent according to $\boldsymbol{p}_i$, $i = 1, ..., n$. Since linear pooling does not preserve independence, independence is not preserved at some of the interior, non-extreme points of $\boldsymbol{P}$. However, the whole set of probability *values for* A, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$, is the same before and after conditionalizing on B.

[8]This binary case of irrelevance can be generalized to non-binary partitions. Let $A_1, ..., A_k$ be a partition of Ω. In Levi's setup, a question is represented as a partion, each element of which is a potential answer. Information B is pairwise irrelevant to $A_1, ..., A_k$ if B is irrelevant to each cell of the partition.

Proposition 5. *Convex IP pooling functions satisfy irrelevance preservation.*

So, while stochastic independence and confirmational irrelevance are equivalent in the precise setting (when $\boldsymbol{p}(B) > 0$), they come apart in the IP context.[9] Because irrelevance preservation reduces to probabilistic independence preservation when the output of the pooling function is a unique probability function, and linear, geometric, and multiplicative pooling functions do not satisfy probabilistic independence preservation in general, we have that linear, geometric, and multiplicative pooling functions do not satisfy irrelevance preservation either. If there are good reasons to require IP pooling functions to satisfy the stronger stochastic independence preservation property, then convex IP pool does not deliver (though there are IP formats that do (Proposition 7)).

Finally, convex IP pooling admits of a simple characterization in terms of the set of *universally admissible means*.[10] We call a function $\mathfrak{m} : [0,1]^n \to [0,1]$ a *mean* on the interval $[0,1]$. We first define a mapping $\mathfrak{M}_n : \mathbb{P}^n \to \mathscr{P}([0,1]^{[0,1]^n})$ by setting for every $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$:

$$\mathfrak{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \Big\{\mathfrak{m} \in [0,1]^{[0,1]^n} : \mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) \in \mathbb{P}\Big\}.$$

Call a mean *admissible* for $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ if $\mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) \in \mathbb{P}$. Then, $\mathfrak{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is the set of admissible means for $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Using $\mathfrak{M}_n$, we define another mapping $\mathcal{M}_n : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$ by setting for every $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$:

$$\mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \Big\{\mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) : \mathfrak{m} \in \bigcap_{\vec{\boldsymbol{q}} \in \mathbb{P}^n} \mathfrak{M}_n(\vec{\boldsymbol{q}})\Big\}.$$

$\mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is the set of probability functions that results from composing each *universally* admissible mean with $\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big)$.

[9] Pedersen and Wheeler show how logically distinct independence concepts are teased apart in the context of imprecise probabilities. They write, "there are several distinct concepts of probabilistic independence and [...] they only become extensionally equivalent within a standard, numerically determinate probability model. This means that some sound principles of reasoning about probabilistic independence within determinate probability models are invalid within imprecise probability models" (2014, p. 1307). So IP provides a more subtle setting in which to investigate independence concepts.

[10] Thanks to Paul Pedersen for suggesting that I include a result along these lines.

Proposition 6. *Suppose that $\mathscr{A}$ contains at least three pairwise incompatible events. A mapping $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$ is a convex IP pooling function—that is, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$—if and only if $\mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ for all $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$.*

An interesting line of research to pursue would be to consider pooling from the perspective of an analysis of means (e.g., Kolmogorov, 1930; de Finetti, 1931; Aczél, 1948). Perhaps such an analysis could shed light on issues like the propriety of qualitative conditions on pooling rules, or the function of convexity in reaching a consensus in inquiry and deliberation.

6. Epistemic and Procedural Grounds for IP Accounts

In their review article, Dietrich and List claim that the question of how probabilities should be aggregated admits of no obvious answer, and, ultimately, the appropriateness of the pooling method depends on the purpose and context of aggregation (2014). They raise the question of whether a pooling method should be justified on *epistemic* or *procedural* grounds. To be justified on epistemic grounds, "the pooling function should generate collective opinions that [...] respect the relevant evidence or track the truth, for example." In order to be justified on procedural grounds, a pooling method should yield a collective opinion that is a "fair representation of the individual opinions" (2014, p. 2). Dietrich and List claim that, while linear pooling can be justified on procedural grounds, it cannot be justified on epistemic grounds. By satisfying WSFP, linear pooling functions reflect "the democratic idea that the collective opinion on any issue should be determined by individual opinions on that issue" (2014, p. 6). Geometric pooling, however, can be justified in epistemic terms "by invoking the axiom of external Bayesianity" (2014, p. 13). The idea seems to be that since updating is a response to the evidence, a pooling method that is well-behaved in the sense of commuting with updating "respects the relevant evidence" by not allowing the order of operations to distort evidential impact.

Convex IP pooling, then, can be justified on Dietrich and List's procedural grounds in virtue of satisfying WSFP. Concerning procedural grounds in general, it is difficult to think of a more fair or democratic representation of individual opinions than a representation that *includes* each opinion and all of the compromises between opinions. But since convex IP pooling functions also satisfy external Bayesianity, it would thus appear that the alleged tension between epistemic and procedural criteria for probabilistic opinion aggregation can be resolved by simply moving to an IP account.

I endorse the basic motivation for Dietirch and List's discussion. Like other deliberate activities, pooling is *goal-directed*. How one should approach pooling depends on, among other things, one's goals. One may have multiple goals, in which case, tensions in jointly satisfying them may require tradeoffs. Nevertheless, I find it rather opaque precisely how WSFP encodes an intuitive procedural constraint on pooling. Similarly, how commutativity of pooling and updating ensures that the collective opinion respects the relevant evidence or tracks the truth stands in need of further clarification. Even if the *philosophical* distinction and interpretation of WSFP and external Bayesianity does not admit of further clarification, however, our point stands that the tension between satisfying the *formal* desiderata can be resolved in the IP setting.

7. Objections to IP Pooling

Perhaps the relative neglect of IP in discussions of pooling can be explained in part by a skepticism concerning the *use* to which IP sets can be put. In their very nice overview of work on pooling, Genest and Zidek write, "the jury remains out on the theory of Walley [...] In particular, it is unclear how [the IP set] could be used 'at the end of the day'" (1986, p. 124). There are essentially two types of uses to which an account of probabilities may be put: those concerning epistemic issues like inference, and those concerning issues in decision making.

7.1. Epistemology. Some degree of the skepticism about the epistemic usefulness of IP may be dispelled by considering recent work. For instance, after Genest and Zidek's article, Walley published his magisterial book addressing applications of IP to issues in statistical reasoning (1991). Fabio Cozman explores the application of IP to issues in Bayesian networks in a number of papers (1998; 2000).

But let us consider some epistemological challenges of a general sort. In reviewing some difficulties for the few available accounts of pooling IP sets of probabilities (accounts allowing imprecision at the individual level), Robert Nau claims that neither taking the union nor the intersection of convex sets of imprecise probabilities yields a satisfactory account of pooling. He writes,

> As more opinions are pooled, the union can only get larger, and it reflects only the least informative opinions, whereas intuitively there ought to be (at least the possibility of) an increase in precision as the pool gets larger. On the other hand, the intersection of convex sets of measures may be empty if experts are mutually incoherent, and it generally yields too tight a representation of aggregate uncertainty. As more opinions are pooled, the intersection can only shrink, and it reflects only the most extreme among those opinions, whereas intuitively there should be some convergence to an average opinon when the pool gets sufficiently large. Moreover, neither the union nor the intersection provides an opportunity for the differential weighting of opinions, which would be desireable in cases where one individual is considered (either by herslef or by an external evaluator) to be better or worse informed than another individual about a particular event under consideration. (2002, p. 267)

Similar concerns could be expressed about the account of pooling under examination in this essay since uncertainty never decreases by mere pooling on our account. But I think they would misplaced. The appropriateness of the behavior of a pooling function cannot be assessed in abstract, without specifying the *point* of pooling probabilities in the first place. If the point is to find common ground among the opinions being pooled, increasing uncertainty is to be expected. In general, the more opinions among which we try to find common ground, the less common ground there will be.[11] One might not wish to seek consensus among certain opinions, but that is a different matter. On our account uncertainty *can* be reduced, but through inquiry and not through pooling. As the

[11] I suspect that Nau is not targeting consensus because his models of pooling involve game-theoretic bargaining scenarios pitting the opinions to be aggregated against each other.

group acquires sufficient information, conditionalization generally leads to a reduction of imprecision. In the IP setting, it is also possible, however, for conditionalization to *increase* imprecision in the short run, a phenomenon known as *dilation* (Seidenfeld and Wasserman, 1993; Wasserman and Seidenfeld, 1994; Herron et al., 1997; Pedersen and Wheeler, 2014, 2015). But our point here is not that conditionalization invariably decreases uncertainty, but that it can and that decreasing uncertainty through conditionalization has familiar Bayesian "learning" foundations whereas pooling (averaging) does not.

Furthermore, in the case of pooling imprecise probabilities, I would not endorse taking intersections for the purpose of finding consensus. In the case of mutual incoherence, intersections yield the empty set. But the lack of any consensus concerning which probability functions can be ruled out means that the group in consensus cannot rule any probability functions out. Taking the convex hull of the union would reflect this, yielding complete uncertainty.[12]

I think it is important to distinguish between finding consensus among some opinions and taking those opinions as evidence. In the latter case, if an agent outside the group considers some members of the group to be less informed than others, *that* opinion should be reflected in conditionalization through the likelihood for the experts' opinions (Cf. the *Supra-Bayesian* approach to pooling (Genest and Zidek, 1986, p. 120)). In the former case, if a group member is considered, by herself or other group members, to be less informed, consensus is often not sought. Finding what common ground the group members share is unproblematic when consensus is sought, regardless of the social, political, or intellectual clout members accord each other. It is also open to, and perhaps rationally obligatory for, the modest group member to allow her opinion concerning her relative informedness to be reflected in her probabilistic opinion before pooling.

Finally, one might object that IP pooling amounts to declining to really aggregate. In a sense, that is true, if pooling is restricted to taking some sort of average of individual probabilities. But,

[12] See Larry Wasserman's review of Walley's book for objections to this representation of complete uncertainty (1993), and Levi's concept of *confirmational commitment* as a potential means of addressing the objections (1974).

again, what is the theoretical basis for only considering precise averages of subjective probabilities? An IP set clearly *represents* group opinion, and can be employed in inference and decision making.

7.2. Decision Theory. Because decision theory is a very involved topic and I do not treat it in this chapter, I limit myself to pointing out that sophisticated decision theories for IP have been developed and extensively studied. These include Levi's E-admissibility and tie-breaking decision rule (1980), Girón and Rios' quasi-Bayesian decision theory (1980), Gilboa and Schmeidler's Γ-Maximin (1989), and Walley's *Maximality* (1991). Seidenfeld, Schervish, and Kadane axiomatize their theory of coherent choice under uncertainty in the framework of set-valued choice functions (2010). Though saying so overcommits me for my project in this essay, I hold the view that theoretical disputes about probability cannot be adjudicated without thorough decision theoretic considerations.

7.3. Convexity Revisited. How much do the results in this essay depend on the convexity of the IP set? Not much! To see why, consider the following very simple IP pooling function, $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$, such that

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \{\boldsymbol{p}_i : i = 1, ..., n\}$$

So defined, $\mathcal{F}$ takes as input a profile of probability functions and returns the set of functions in that profile.

Proposition 7. *Let $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$ be an IP pooling function such that, for each profile in $\mathbb{P}^n$, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \{\boldsymbol{p}_i : i = 1, ..., n\}$. Then, $\mathcal{F}$ satisfies SSFP, WSFP, ZPP, MP, unanimity preservation, external Bayesianity, and confirmational irrelevance preservation. Moreover, $\mathcal{F}$ satisfies stochastic independence preservation.*

The proof of Proposition 7 is straightforward and so is omitted here. The upshot is that pooling with imprecise probabilities is promising in a robust sense. So, while the convex IP pooling model is the chief subject of our philosophical approbation and mathematical analysis in this essay, our case for the consideration of IP in the context of pooling does not rest exclusively with that model.

7.4. Dynamics. As I have presented convex IP pooling functions, the input is a profile of individual probability functions and the output is a convex set of probability functions. What if individual probabilities are themselves imprecise? Or what happens if we attempt to pool the probabilistic opinions of agents that are themselves groups?[13] As it stands, our account is silent. There is, however, a natural extension of the account on offer. Consonant with the philosophical position staked out here, the idea is to convexify the *union* of sets of probability functions.

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

Where the profile consists of n *sets* of probability functions, $(\boldsymbol{P}_1, ..., \boldsymbol{P}_n)$, the pool is given by $\mathcal{F}(\boldsymbol{P}_1, ..., \boldsymbol{P}_n) = conv\{\bigcup_i \boldsymbol{P}_i\}$. I leave examination of this more complete account to future work.

8. Conclusion

According to standard Bayesian theory, personal probabilities are *subjective*. One route that has been explored for recovering some objectivity is establishing intersubjective agreement. There are, for example, the famous convergence theorems to the effect that, given non-extreme priors and a suitably large amount of evidence upon which to conditionalize, posteriors converge (Savage, 1954; Gaifman and Snir, 1982). Consensus in the (*very*) long-run, however, is not the only kind of consensus we may seek. Prior to inquiry, consensus as *shared agreement* is still possible, and desirable for joint efforts in inquiry. Convex IP pooling can be philosophically motivated as an account of such consensus.

[13] The problem being raised is similar to one in the literature on AGM belief revision. The *principle of categorical matching* requires that the output of a belief revision operator be of the same format as the input. Otherwise, the account of belief revision, constructed for a certain input format, is silent about iterated belief revision (Gärdenfors and Rott, 1995). In the case of convex IP pooling functions, dynamics *of pooling* are defined so long as we are never pooling sets of probabilities.

Our objective has been to undermine the preconception that probabilistic opinion pooling should result in a representative probability function for the group. Our tack has been to explore another option, arguing that, even by the very lights of those working in pooling, this option is promising. We have the following summary (an “X” means the pooling method does not generally satisfy the property):

Table 1. Pooling Method Report Card

	Linear	**Geometric**	**Multiplicative**	**Convex IP**
SSFP	✓	X	X	✓
ZPP	✓	✓	✓	✓
MP	✓	X	X	✓
WSFP	✓	X	X	✓
Unanimity Preservation	✓	✓	X	✓
External Bayesianity	X	✓	X	✓
Individualwise Bayesianity	X	X	✓	X
Irrelevance Preservation	X	X	X	✓

Perhaps the most sensible representation of group opinion, especially when pooling is interpreted as reaching consensus, is not in terms of a single probability function. At the very least, the arguments and results above may be read both as an exploration of extending the mathematical framework of opinion pooling to cover IP pooling, and as a plea for liberalism about pooling formats.

CHAPTER 3

Learning and Pooling, Pooling and Learning

1. Introduction

Bayesian conditionalization is the gold standard of probabilistic learning. Yet several authors advocate modifications of conditionalization for a number of reasons. For example, conditionalization entails assigning probability 1 to the evidence. Dissatisfied with such "dogmatic epistemology," Richard Jeffrey proposed his *probability kinematics* as a way of updating on *uncertain* evidence (Jeffrey, 2004). To take another example, consider probabilistic *imaging*. One widely pursued goal in work on the logic of conditionals is to find a way of identifying the probability of a conditional with the corresponding conditional probability. Attempts to do so have been repeatedly frustrated by a series of triviality results. However, David Lewis introduces imaging and shows that a version of the identity holds when formulated in terms of imaging instead of in terms of conditionalization (Lewis, 1976). And there are other proposals, such as minimizing the Kullback-Leibler divergence (Kullback and Leibler, 1951). Here, the objective is to accommodate the evidence in such a way as to minimize a measure (the K-L divergence) of the difference between posterior and prior.

Probabilistic opinion pooling can be viewed as part of an important strand in Bayesian epistemology and statistics concerned with consensus. The received view is that personal probabilities are *subjective* (Ramsey, 1990b; Savage, 1954; de Finetti, 1964), resulting in much fretting about the implications for scientific methodology. The worry is that the objectivity of scientific confirmation, explanation, inference, and the like is compromised to the extent that such probability plays a role.

A version of this chapter is published as a paper, also coauthored with Ignacio Ojea Quintana, under the same title in *Erkenntnis*.

A prominent Bayesian response comes in the form of convergence and merging of opinions theorems, which show that, given agreement about probability 0 events and enough evidence, probabilities converge (almost surely) (Savage, 1954; Gaifman and Snir, 1982).[1] Conditionalization, that is, leads to consensus, "washing out" the problematically subjective priors leaving intersubjective agreement in their place.

Pooling offers a distinct way of reaching a consensus in probabilistic opinion, one available even when the opportunity to collect more evidence is not. Consensus is reached immediately *via* methods for *aggregating* probabilistic judgments instead of in the long run *via* conditionalization (Huttegger, 2015b). After all, as Keynes astutely observes, we have reasons not to be particularly concerned with the long run. As with convergence arguments, intersubjective agreement stands in for objectivity in the pooling context. Still, in the literature on probabilistic opinion pooling, one of the constraints of central concern is *external Bayesianity*, which requires that pooling and Bayesian conditionalization on a common likelihood function commute (Madansky, 1964). That is, the result of pooling and then updating is the same as first updating and then pooling. The order of operations does not change the outcome. One natural question to ask in light of the aforementioned alternatives to conditionalization is, what about commutativity with alternative updating policies? This is the question that concerns us in the present essay.

Much of the focus in the pooling literature is on characterization and impossibility results. Such results are not the intended contribution of this chapter (though Proposition 8 generalizes a characterization result due to Wagner to the imprecise probabilities setting). Instead, we continue an exploration of the potential of imprecise probabilities in the context of learning and pooling. In the previous chapter, I argued that collective opinion is more properly represented by imprecise probabilities (IP) in general. I provide three arguments. First, if pooling is interpreted as reaching a *consensus* in probabilistic opinion, IP pooling is on firmer philosophical ground (Cf. Levi, 1985;

[1] Not all merging of opinions results require probabilities to converge to certainty (Blackwell and Dubins, 1962). Under certain conditions, Bayesian conditionalizing can bring probabilities close even if they do not converge to 1 or 0.

Seidenfeld et al., 1989). The point, briefly, is that IP models allow for suspending judgment between some number of probability distributions by not ruling them out for use in deliberation and inquiry, and reflect the common ground among the group concerning which probability distributions *are* ruled out. Such a consensus constitutes a neutral initial position from which to launch further inquiry. Precise pooling functions, on the other hand, do not allow for an analogous suspense of judgment, and may yield collective probabilistic opinions endorsed by none of the group members. Second, there are IP pooling functions that jointly satisfy more of the standard pooling constraints than any precise pooling recipe can. Third, in the IP setting, the tension between a pooling method's being justified on epistemic or procedural grounds (Dietrich and List, 2014)—reflected in the tension between satisfying certain formal epistemic and procedural constraints—dissipates, an artifact of the assumption of precision.

The results that follow may be taken to contribute to that case for IP. Briefly put, I show that, while the form of updating for a given precise pooling method is quite restricted under the requirement of commutativity, relaxing the assumption that the collective opinion should take the form of a numerically determinate probability function enables us to lift many of those restrictions. Several revision methods are consistent with pooling understood the IP way. After introducing the mathematical pooling framework in the next section, we begin with the gold standard in Section 3. I remain neutral as to whether Jeffrey conditionalization and imaging ultimately admit of sufficient motivation, though I rehearse some of the standard motivations for and reservations about probability kinematics (Section 5) and imaging (Section 6), and state the commutativity results. Motivations for requiring commutativity of learning and pooling are discussed Section 4.

2. Preliminaries

Let Ω denote a sample space, a set of mutually exclusive and exhaustive possible states of the world.[2] In what follows, we assume that Ω is countable. A function $\boldsymbol{p} : \Omega \to [0, 1]$ is a *probability mass function* (pmf) *iff* $\sum_{\omega \in \Omega} \boldsymbol{p}(\omega) = 1$. An algebra $\mathscr{A}$ of events over Ω is a set of subsets of Ω closed under complementation and finite unions; closure under countable unions yields a σ-algebra. We assume throughout the essay that $\mathscr{A}$ is a σ-algebra. Given a pmf, we can define a *probability measure*, (abusing notation by using the same symbol as for pmfs) $\boldsymbol{p}$, on events in general: $\boldsymbol{p}(E) = \sum_{\omega \in E} \boldsymbol{p}(\omega)$.

Let $\mathbb{P}$ be the set of all pmfs on Ω. A *precise* pooling function is a function, $F : \mathbb{P}^n \to \mathbb{P}$, mapping a profile of n pmfs, $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, to a single pmf, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Typically, the n pmfs are taken to represent the opinions of a set N of individuals, and $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is supposed to represent, in some sense, the aggregate or collective opinion. Various candidate interpretations of an opinion pool exist in the literature: a rational consensus (adopted genuinely by all members or adopted merely "for the sake of the argument"); 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. (Genest and Zidek, 1986; Wagner, 2009). There are a number of concrete pooling functions discussed in the literature, but, by far, the two most prominent are linear and geometric pooling functions.

Linear Opinion Pools. $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i$, where $\alpha_i \geqslant 0$ and $\sum_{i=1}^{n} \alpha_i = 1$.

A linear opinion pool is just a weighted arithmetic average of the n probability functions. A geometric pooling function takes the (weighted) *geometric* average of the n pmfs.

[2] Ω may be thought of as a partition of a space of agent-relative serious possibilities determined by consistency with a state of full belief. As is a state of full belief, Ω is open to being revised, refined, etc., as judged appropriate (Levi, 1980).

Geometric Opinion Pools. $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = c\prod_{i=1}^{n} \boldsymbol{p}_i^{\alpha_i}$, where $\alpha_i \geqslant 0$ and $\sum_{i=1}^{n} \alpha_i = 1$, and $c = \frac{1}{\sum_{\omega' \in \Omega}[\boldsymbol{p}_1(\omega')]^{\alpha_1} \cdots [\boldsymbol{p}_n(\omega')]^{\alpha_n}}$ is a normalization factor.[3]

The focus of this section of the dissertation is on a generalization of pooling functions to the IP setting: $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$.[4] We use $\mathcal{F}$, opposed to F, to denote set-valued pooling operators. IP pooling functions are maps from profiles of probability measures to sets of probability measures. In particular, we consider pooling functions that map profiles of n pmfs to the convex hull of those functions: $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv\{\boldsymbol{p}_i : i = 1, ..., n\}$. The convex hull is the smallest convex set containing $\boldsymbol{p}_i$ for $i = 1, ..., n$. A set, $\boldsymbol{P}$, is convex if it satisfies the following property.

Convexity. If $\boldsymbol{p}_1, \boldsymbol{p}_2 \in \boldsymbol{P}$, then $\alpha\boldsymbol{p}_1 + (1 - \alpha)\boldsymbol{p}_2 \in \boldsymbol{P}$ for $\alpha \in [0, 1]$. [5]

Put another way, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is the set of all convex combinations of the individual probability functions. We let $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$ denote the set of probability values assigned to A:

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$$

There are alternative IP formats, including set-based formats that do not require convexity, interval-valued probability functions, or, more operationally, sets of desirable gambles, for instance. In my view, all are worthy of extensive study. Convex sets are quite commonly employed in the theory of IP, including in sophisticated decision theories. Since I do not intend to settle complex debates internal to IP theory here, and the results to come do not hinge on convexity in the sense that there are alternative IP formats for which they hold, I will simply restrict my attention to convex sets as an illustration of the potential of IP in theorizing about pooling.

[3] Notice that, due to the way geometric pooling is defined, there are profiles for which $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(\omega) = 0$ for all $\omega \in \Omega$—in violation of the probability axioms. Such a situation arises if for each $\omega \in \Omega$ there is a $\boldsymbol{p}_i \in (\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ such that $\boldsymbol{p}_i(\omega) = 0$. Circumventing this problem, Wagner restricts the domain of pooling operators to the set of profiles for which this does not happen. That is, the domain of a pooling function is the set of profiles such that there is some $\omega \in \Omega$ for which $\boldsymbol{p}_i(\omega) > 0$ for all $i = 1, ..., n$.

[4] See (Schervish and Seidenfeld, 1990; Herron et al., 1997) for studies of convergence relevant to IP.

[5] Within the IP research community, convexity is a matter of some controversy. For attacks on the requirement, see (Seidenfeld et al., 1989; Kyburg and Pittarelli, 1992; Seidenfeld et al., 2010). For defenses, see (Levi, 1990, 2009).

3. External Bayesianity

Essential to Bayesian methodology is the assumption of a *prior* probability distribution on the algebra of events (or propositions) of concern. Learning proceeds by *conditionalizing* the prior on the evidence, yielding a *posterior* distribution. Conditionalization of a probability function, $\boldsymbol{p}$, on evidence, E, results from setting the posterior probability for any event $A \in \mathscr{A}$ equal to the prior conditional probability $\boldsymbol{p}(A|E)$.

$$\boldsymbol{p}^E(A) = \boldsymbol{p}(A|E) = \frac{\boldsymbol{p}(A \cap E)}{\boldsymbol{p}(E)}, \text{ when } \boldsymbol{p}(E) > 0.$$

The posterior, $\boldsymbol{p}^E$, can be thought of as the result of learning E. In the context of sets of probability functions, conditionalization can be generalized by conditionalizing each member of the set:

$$\mathcal{F}^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \{\boldsymbol{p}^E : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n),\ \boldsymbol{p}(E) > 0,\ \text{and } \boldsymbol{p}^E(\cdot) = \boldsymbol{p}(\cdot|E)\}$$

Call $\mathcal{F}^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ the *prior-by-prior* conditionalization of $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ (when $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is convex, $\mathcal{F}^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is called the *convex* conditionalization of $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$).[6] We define prior-by-prior conditionalization generally, allowing $\boldsymbol{p}(E) = 0$ for some $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ (Cf. Kyburg, 1987, p. 279). But when we first update the $\boldsymbol{p}_i$, we assume $\boldsymbol{p}_i(E) > 0$ for $i = 1, ..., n$; otherwise, $\mathcal{F}(\boldsymbol{p}_i^E, ..., \boldsymbol{p}_n^E)$ is not defined. Though not stated in the language of probabilistic opinion pooling, proofs of the commutativity of convexifying a set of probability functions and conditionalization exist in the literature.

THEOREM 6. *(Levi, 1978; Girón and Ríos, 1980) Convex IP pooling commutes with conditionalization.*

[6] In the IP setting, conditionalization can actually lead to *greater* uncertainty in the short-run, a very interesting phenomenon known as *dilation* (Seidenfeld and Wasserman, 1993; Pedersen and Wheeler, 2014).

Figure 1. Commutativity of Pooling and Conditionalization

$$
\begin{array}{ccc}
 & (\boldsymbol{p}_1, ..., \boldsymbol{p}_n) & \\
 \swarrow & & \searrow \\
\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) & & (\boldsymbol{p}_1^E, ..., \boldsymbol{p}_n^E) \\
\downarrow & & \downarrow \\
\mathcal{F}^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) & \longleftrightarrow & \mathcal{F}(\boldsymbol{p}_1^E, ..., \boldsymbol{p}_n^E)
\end{array}
$$

Importantly, linear opinion pooling does not commute with conditionalization, though geometric pooling does (Genest, 1984; Russell et al., 2015). As we will see, linear pooling does commute with imaging, though geometric pooling does not.

As standardly pointed out, *external Bayesianity* is a generalization of the requirement that pooling and standard conditionalization commute (Wagner, 2009; Dietrich and List, 2014; Russell et al., 2015), because conditionalization on a common likelihood function generalizes standard Bayesian conditionalization on an event. A *likelihood function*, $\lambda : \Omega \to [0, \infty)$, is intended to encode, given any $\omega \in \Omega$, how expected some evidence is with the number $\lambda(\omega)$. The conditionalization of a pmf, $\boldsymbol{p}$, on a likelihood function, λ, is give by the following formula.

$$\boldsymbol{p}^{\lambda}(\omega) = \frac{\boldsymbol{p}(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')}, \text{ when } \sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega') > 0$$

For the special case of Bayesian conditionalization on an event E, define λ as the indicator function for that event:

$$\lambda(\omega) = \begin{cases} 1, \text{ if } \omega \in E \\ 0, \text{ otherwise.} \end{cases}$$

External Bayesianity requires that updating the individual probabilities on a common likelihood function and then pooling is the same as pooling and then updating the pool on that likelihood function.

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

When pooling is presumed to produce a numerically determinate probability function for the group, generalized geometric pooling functions uniquely satisfy external Bayesianity (Genest et al., 1986; Nau, 2002). Extended to the IP setting, the constraint requires $\mathcal{F}(\boldsymbol{p}_1^{\lambda}, ..., \boldsymbol{p}_n^{\lambda}) = \mathcal{F}^{\lambda}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. The requirement is still that learning by updating on a common likelihood function and pooling commute, but the format of the pool is altered. What is not altered, we submit, are the compelling aspects of the constraint. Recall the following observation from the previous chapter.

PROPOSITION 3. Convex IP pooling functions are externally Bayesian.

(The *propositions* in this dissertation are my results. When provided, proofs are relegated to the appendix.) And since Bayesian conditionalization is a special case of updating on a likelihood function, it follows that any IP pooling function (not necessarily convex) that is externally Bayesian also satisfies commutativity with conditionalization.

The fact that updating on a common likelihood generalizes updating on an event may serve to show that the assumption of a common likelihood function is not quite as strong as it may appear initially, since the conditionalization of the $\boldsymbol{p}_i$ on some event drops out as a special case. That is, learning the same event is an instance of a shared likelihood function. It is also worth pausing to consider why Bayesians would deal in likelihood functions in the first place if updating with a likelihood function presents ways of learning not reducible to conditionalization.[7] One reason is that, under certain conditions, there is a way of regarding updating with a likelihood function as a case of Bayesian conditionalization by *refining* the algebra so that there is an event such that conditionalizing on it yields the same results as updating with the likelihood function on the coarser algebra. We return to this point—which is relevant to Jeffrey Conditionalization as well—at the close of Section 5.

[7]Thanks to Paul Pedersen for emphasizing this point to me.

4. Commutativity

But why should it matter if a pooling method is externally Bayesian? More generally, why should we insist on the commutativity of learning and pooling? A few motivations, which I now briefly survey, are offered in the literature.

In introducing the external Bayesianity constraint, Madansky points out that the decisions of a group with common interests employing an externally Bayesian pooling operator will appear to outsiders as the decisions of a single Bayesian agent (1964). How? A Bayesian agent conditionalizes. So, given group opinion, $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, the updated group opinion should result from the group prior by conditionalization, $F^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. If the group employs a pooling operator that is not externally Bayesian, and learning happens at the level of individuals, the posterior group opinion may not result from the prior group opinion by conditionalization: $F(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) \neq F^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. If the relevant learning happens at the level of group opinion, then the posterior group opinion may not be the result of applying the (non-externally Bayesian) pooling method that allegedly gives us the way of arriving at group opinion when applied to individual opinions: again, $F(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) \neq F^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$.

Of course, there are a number of interpretations of the pool of individual opinions, including as "a rough summary" of the n pmfs (Wagner, 2009). The properties that are appropriate for a pooling function to exhibit depend on the interpretation of the pool and the use to which it is put. When pooling is interpreted as a method of reaching either a compromise or a genuine consensus for use in group decision making, it may be important to ensure adherence of "group opinion" to norms of individual rationality. For views according to which groups can be agents subject to the same rationality conditions, for example, failing to satisfy external Bayesianity raises problems insofar as Bayesian conditionalization is rationally mandatory. More generally, those problems arise for failures of commutativity of pooling with any rule of learning held to be normatively compelling.

Furthermore, Russell et al. charge pooling methods that fail to commute with conditionalization with vulnerability to a diachronic Dutch book (2015). When the group posterior does not result from the group prior by conditionalizing on what the individuals learn, the conditions for a diachronic Dutch book (at the level of group opinion) are met. Echoing Raiffa (1968, pp. 221-226), Dietrich and List offer other strategic considerations in favor of external Bayesianity. If a pooling method is not externally Bayesian, collective opinion is open to manipulation. By disclosing relevant information at the appropriate time, someone could affect collective opinion by increasing the influence of certain opinions, for example (2014). There are, in short, possible manipulations besides those of a clever bookie.

Perhaps most uncontroversially, pooling operators for which learning and pooling commute save us the trouble of having to figure out whether updating should come before or after pooling, whether susceptibility to a diachronic book is damning for the pooling method, how and when to safeguard against manipulation, etc. In any event, the main position argued for in this chapter can be understood as a conditional: *if* one finds commutativity of learning (of various types) and pooling compelling, *then* one has reasons to seriously consider IP pooling formats.

5. Jeffrey Conditionalization

As indicated in the introduction, standard Bayesian conditionalization requires that the event conditionalized upon receives probability 1 in the posterior distribution. Jeffrey's point is that not all learning experiences are like that. Sometimes observation leads to a revision in subjective probability even when there is no proposition (event) E that is learned "for certain." Jeffrey's famous candle light example serves to illustrate his point. Suppose you observe your friend's coat, but only under candle light. The coat looks blue, but you are not quite sure. The impact of this observation is a shift in your subjective probabilities concerning the color of the coat, but none of the options goes to 1. This sort of scenario, some Jeffrey sympathizers claim, is "the normal case" (e.g., Spohn, 2012, p. 38). Improved lighting generally only shifts probabilities a bit more.

Let $\boldsymbol{E} = \{E_i\}$ be a countable family of pairwise disjoint events partitioning Ω. In the candle light example above, the partition of concern consists of the possible colors of the coat. A posterior, $\boldsymbol{q}$, comes from a prior, $\boldsymbol{p}$, by Jeffrey conditionalization by updating on the new probabilities for the cells of $\boldsymbol{E}, \boldsymbol{q}(E_i)$, in the following way:

$$\boldsymbol{q}(A) = \sum_i \boldsymbol{q}(E_i)\boldsymbol{p}(A|E_i)$$

The $\boldsymbol{q}(E_i)$ express the *direct* effect of an observation on subjective probabilities for the cells of the partition. When $\boldsymbol{q}(E_i) = 1$ for some E_i, Jeffrey conditionalization reduces to standard Bayesian conditionalization.

A fact about Jeffrey conditionalization that many have found problematic is the failure of learning sequences to commute (Skyrms, 1986).

FIGURE 2. Commutativity of Updatings

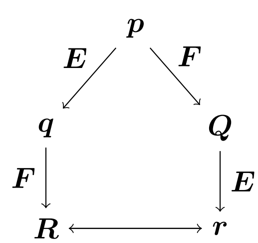

Suppose that the learning experiences prompting the revision from $\boldsymbol{p}$ to $\boldsymbol{q}$ and the revision of $\boldsymbol{Q}$ to $\boldsymbol{r}$ are the same, as are those leading to the revision of $\boldsymbol{q}$ to $\boldsymbol{R}$ and $\boldsymbol{p}$ to $\boldsymbol{Q}$ (reflected in the updating partitions $\boldsymbol{E}$ and $\boldsymbol{F}$, respectively). Jeffrey conditionalization can yield $\boldsymbol{R} \neq \boldsymbol{r}$. That is, switching the order of two learning experiences can yield different probabilities in the end. Van Fraassen complains:

> Two persons, who have the same relevant experiences on the same day, but in a different order, will not agree in the evening even if they had exactly the same opinions in the morning. Does this not make nonsense of the idea of learning from experience? (1989, p. 338)

It is because of the issue of commutativity of learning experiences (as well as the nice off-the-shelf result of Wagner's (Theorem 8) we appeal to below) that we present here a particular parameterization of Jeffrey conditionalization intended to address the commutativity difficulty.

Hartry Field offers a fix, identifying conditions that are sufficient to ensure that, for finite partitions, $\boldsymbol{R} = \boldsymbol{r}$ in Figure 2 above (Field, 1978). Wagner generalizes the result to countable partitions (Wagner, 2002). We introduce some useful notation. Where A and B are events and $\boldsymbol{q}$ is a revision of $\boldsymbol{p}$, the *Bayes factor* is the ratio of new to old odds:

$$\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B) = \frac{\boldsymbol{q}(A)/\boldsymbol{q}(B)}{\boldsymbol{p}(A)/\boldsymbol{p}(B)}$$

Instead of being reflected in identical *posteriors*, the proposal on the table is to understand identical learning as reflected in identical Bayes factors. Wagner points out that identifying identical learning with identical Bayes factors has a distinguished pedigree in Bayesian thinking (Good, 1983; Jeffrey, 2004).[8] What Field shows is that Jeffrey conditionalization is commutative when *identical learning* is interpreted as identical Bayes factors.

[8] Wagner contends that identical learning should be thought of as identical Bayes factors rather than identical posteriors. One alleged reason is that posteriors are tainted by the prior, whereas Bayes factors are an uncontaminated measure of the impact of the evidence. How do Bayes factors measure the impact of the evidence in isolation from the prior? Consider the case in which $\boldsymbol{q}$ comes from $\boldsymbol{p}$ by Bayesian conditionalization on E. Then,

$$\boldsymbol{q}(A)/\boldsymbol{q}(B) = \frac{\boldsymbol{p}(A|E)}{\boldsymbol{p}(B|E)}$$

and

$$\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B) = \frac{\boldsymbol{p}(A|E)/\boldsymbol{p}(B|E)}{\boldsymbol{p}(A)/\boldsymbol{p}(B)}.$$

So, $\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B)$ is a measure of the change the evidence, E, induces in favor of A over B. $\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B)$ can also be rearranged using Bayes' theorem.

$$\frac{\boldsymbol{q}(A)}{\boldsymbol{q}(B)} = \frac{\boldsymbol{p}(A|E)}{\boldsymbol{p}(B|E)} = \frac{\frac{\boldsymbol{p}(A)\boldsymbol{p}(E|A)}{\boldsymbol{p}(E)}}{\frac{\boldsymbol{p}(B)\boldsymbol{p}(E|B)}{\boldsymbol{p}(E)}} = \frac{\boldsymbol{p}(A)\boldsymbol{p}(E|A)}{\boldsymbol{p}(B)\boldsymbol{p}(E|B)} = \frac{\boldsymbol{p}(A)}{\boldsymbol{p}(B)} \times \frac{\boldsymbol{p}(E|A)}{\boldsymbol{p}(E|B)}$$

Dividing now by $\frac{\boldsymbol{p}(A)}{\boldsymbol{p}(B)}$, the denominator of $\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B)$, gives us

$$\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; A : B) = \frac{\boldsymbol{p}(E|A)}{\boldsymbol{p}(E|B)}$$

The quantity $\boldsymbol{p}(E|A)/\boldsymbol{p}(E|B)$ is sometimes referred to as the *likelihood ratio*. So, the Bayes factor is a ratio of the non-prior quantities involved in Bayes' theorem, the quantities that revise the prior.

Theorem 7. *(Wagner, 2002, Theorem 3.1) Consider the revision scheme of Figure 2. If*

$$\mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; E_{i_1} : E_{i_2}) = \mathcal{B}(\boldsymbol{r}, \boldsymbol{Q}; E_{i_1} : E_{i_2}) \text{ for all } i_1, i_2$$

and

$$\mathcal{B}(\boldsymbol{R}, \boldsymbol{q}; F_{j_1} : F_{j_2}) = \mathcal{B}(\boldsymbol{Q}, \boldsymbol{p}; F_{j_1} : F_{j_2}) \text{ for all } j_1, j_2,$$

then $\boldsymbol{R} = \boldsymbol{r}$.

Wagner further shows that Jeffrey conditionalization has an equivalent parameterization in terms of Bayes factors (2009, Theorem 3.2). The function $\boldsymbol{q}$ gives us posteriors for atomic events, $b_k = \mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; E_k : E_1)$, $k = 1, 2, ...$, and $[\omega \in E_k]$ is the characteristic function of the set E_k:

$$[\omega \in E_k] = \begin{cases} 1, & \text{if } \omega \in E_k \\ 0, & \text{otherwise.} \end{cases}$$

Wagner's parameterization, then, is the following.

$$\boldsymbol{q}(\omega) = \boldsymbol{p}_J^{\boldsymbol{E}}(\omega) = \frac{\sum_k b_k \boldsymbol{p}(\omega)[\omega \in E_k]}{\sum_k b_k \boldsymbol{p}(E_k)}$$

There are two very nice features of this parameterization that are relevant. First, as I have been explaining, it responds to the complaints about commutativity because the result of a sequence of updates is invariant under permutations of that sequence when Jeffrey conditionalization is understood this way, with identical learning reflected in identical Bayes factors instead of identical posteriors.

The second nice feature, as Wagner shows, is that with his parameterization, we can articulate a version of commutativity with Jeffrey conditionalization that provides us with a characterization of externally Bayesian pooling operators in terms familiar to formal epistemologists and philosophers of science. We call Wagner's version of commutativity with Jeffrey CJC_W.

$\boldsymbol{CJC_W}$. For all partitions $\boldsymbol{E} = \{E_k\}$ of Ω, all profiles $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ in the domain of F, the Jeffrey update of the pool, $F_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \frac{\sum_k b_k F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)[\cdot \in E_k]}{\sum_k b_k F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(E_k)}$, is identical to $F(\frac{\sum_k b_k \boldsymbol{p}_1[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}_1(E_k)}, ..., \frac{\sum_k b_k \boldsymbol{p}_n[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}_n(E_k)}) = F(\boldsymbol{p}_{1J}^{\boldsymbol{E}}, ..., \boldsymbol{p}_{nJ}^{\boldsymbol{E}})$, the pool of the (Jeffrey updated) posteriors.[9]

Crucially, the Bayes factors, b_k for $k = 1, 2, ...$, are held fixed across the $\boldsymbol{p}_i$ (and also used in updating $F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$). This is quite different from holding fixed a common posterior distribution, $\boldsymbol{q}$, in Jeffrey conditionalizing the $\boldsymbol{p}_i$. Put another way, he shows that External Bayesianity is equivalent to CJC_W.

THEOREM 8. *(Wagner, 2009, Theorem 3.3) A (precise) pooling operator is externally Bayesian iff it satisfies CJC_W.*

We take $\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ to be given by Jeffrey conditionalization of each element of $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ on the partition $\boldsymbol{E}$, holding fixed the Bayes factors b_k for $k = 1, 2, ...$ for $\boldsymbol{p}_i$ for $i = 1, ..., n$. A very slight mathematical generalization allows us to extend Wagner's result to IP pooling functions in general.

PROPOSITION 8. *Let $\mathcal{F} : \mathbb{P}^n \to \mathscr{P}(\mathbb{P})$ be an IP pooling function (not necessarily convex). $\mathcal{F}$ is externally Bayesian iff $\mathcal{F}$ satisfies CJC_W.*

In particular, moving to the convex IP setting does not break the equivalence between external Bayesianity and commutativity with Wagner's parameterization of Jeffrey conditionalization. Putting Propositions 3 and 8 together, we obtain the following.

PROPOSITION 9. *Convex IP pooling satisfies CJC_W.*

[9] Wagner's version of commutativity with Jeffrey conditionalization involves some additional technical assumptions. First, that $\boldsymbol{p}_i(E_k) > 0$ for all i and all k. Second, that $b_1 = 1$ and $\sum_k b_k \boldsymbol{p}_i(E_k) < \infty$ for $i = 1, ..., n$. Third, where $\boldsymbol{q}_i(\omega) = \frac{\sum_k b_k \boldsymbol{p}_i(\omega)[\omega \in E_k]}{\sum_k b_k \boldsymbol{p}_i(E_k)}$, it is the case that $0 < \sum_k b_k F(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(E_k) < \infty$. In the IP setting, this last assumption may be adjusted to be a requirement for each $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$.

We have the following immediate corollary (the proof is trivial given Proposition 9, and is omitted).

PROPOSITION 10. *Convexity is preserved under Jeffrey conditionalization with common Bayes factors.*

To say that convexity is preserved means that, if we start with a convex set, we do not lose convexity in moving to the set of updated probability functions.

However, when $\boldsymbol{q}_i$ comes from $\boldsymbol{p}_i$ by standard Jeffrey conditionalization on some shared posterior distribution, $\boldsymbol{q}$, over a partition, $\boldsymbol{E}$, and the pool is updated likewise by updating each element of $\mathcal{F}(\boldsymbol{p}_i, ..., \boldsymbol{p}_n)$ on that same posterior distribution over $\boldsymbol{E}$, Jeffrey conditionalization and convex IP pooling do *not* commute.

PROPOSITION 11. *Convex IP pooling does* ***not*** *commute with Jeffrey conditionalization on a common posterior.*

Instead of holding fixed common Bayes factors, here we assume a fixed posterior distribution on $\boldsymbol{E}$. Neither is commutativity of Jeffrey conditionalization so formulated and pooling guaranteed in the precise setting (Wagner, 2009, pp. 340-341). In particular, linear and geometric pooling fail to commute with Jeffrey in general. Furthermore, certain "objective" Bayesian updating methods, like minimizing the Kullback-Leibler divergence between posterior and prior, generalize Jeffrey conditionalization (Diaconis and Zabell, 1982). A corollary of Proposition 11, then, is that minimizing the Kullback-Leibler divergence does not commute with convex IP pooling. Minimizing the Kullback-Leibler divergence is also a generalization of Jaynes' Maximum Entropy formalism (e.g., Williams, 1980). While there are many advocates of the Kullback-Leibler approach (e.g., Hartmann, 2014), even in the precise setting, a number of Bayesian-style objections to MaxEnt methods have been voiced in the literature (see, e.g., Seidenfeld, 1986; Gaifman and Vasudevan, 2012).

How much ground does Jeffrey conditionalization ultimately gain over standard Bayesian conditionalization? In certain cases—characterized by the *superconditioning criterion*—Jeffrey conditionalization can be represented as Bayesian updating on *certain* evidence in an enlarged probability space (Diaconis and Zabell, 1982, Theorem 2.1). For finite algebras, the effect of Jeffrey conditionalization can always be so represented by Bayesian conditionalization (fn. 1, Wagner, 2002, p. 268). Kyburg attributes an early, informal version of this result to Levi (Levi, 1967b; Kyburg, 1987, Lemma A.5). The point, as Kyburg puts it, is not that we make effable the ineffable observational input motivating Jeffrey conditionalization: "This is not to say that we need to *specify* that evidence; it is that there is an algorithm by means of which the impact of the uncertain evidence can be represented as the impact of other 'certain' evidence" (Kyburg, 1987, p. 280). The obvious question is whether or when the formal superconditioning move is a philosophically legitimate one.

6. Imaging

An hypothesis occupying the attention of many scholars working on the logic of conditionals asserts that the probability of a conditional is identical to the relevant conditional probability: $\boldsymbol{p}(A \to B) = \boldsymbol{p}(B|A)$. There are, of course, a number of ways to interpret the components of the claim. There are, for instance, various ways to define a conditional probability, $\boldsymbol{p}(B|A)$, just as there are various ways to specify an interpretation of a conditional connective, $\to$. Suppose we take conditional probability to be defined standardly as in Section 3. Is there an interpretation of $\to$ such that the desired identity holds for all $\boldsymbol{p}$? No, on pain of triviality, as Lewis proved (1976). That is, a conditional satisfying the identity exists only for trivial probability models. Similar triviality results hold for alternative readings of the identity. For example, "for any $\boldsymbol{p}$ there exists some $\to$ such that $\boldsymbol{p}(A \to B) = \boldsymbol{p}(B|A)$" runs into similar problems. An impressive battery of such triviality results has been obtained for different ways of reading the identity. A helpful overview of much of the relevant literature can be found in (Hájek and Hall, 1994). Though it fails in general when formulated in terms of conditionalization, Lewis shows that a version of the identity holds if formulated in terms of *imaging* instead. We turn now to a brief presentation of imaging and Lewis' possibility result.

Robert Stalnaker specified the semantics of the so-called *Stalnaker conditional*, $>$, in terms of possible worlds. Ω is interpreted as a set of possible worlds.[10] Propositions are subsets of Ω. We assume that $\mathscr{A}$ is a σ-algebra of subsets of Ω, the set of relevant propositions. For any $\omega \in \Omega$, let ω_A be the "most similar" possible world at which A holds, the "closest" A-world to ω. Say that $A > B$ is true at ω *iff* B is true at ω_A (when the antecedent is impossible, $A > B$ is taken to be vacuously true at all worlds). Lewis tailors a probabilistic revision scheme to the Stalnaker semantics.

For any non-empty $E \in \mathscr{A}$, imaging shifts the probability from each $\omega' \in \Omega$ to its "image" atom, $\omega \in E$. If $\omega' \in E$, then ω' is its own image atom. Lewis offers an interpretation in terms of possible worlds. On the assumption that for each world there is a unique "most similar" world in E, imaging can be thought of as the process of revising probabilities by shifting the total probability of each world to its most similar world in E. Relaxing the uniqueness assumption results in what is known as *general imaging*. General imaging allows the probability of each $\omega' \in \Omega$ to be shifted to an image *set*, each element of which receives some fraction of the total probability of ω'. Clearly, general imaging reduces to imaging when the image set is a singleton and the total probability of each $\omega' \in \Omega$ is transferred to its image set.

Formally, we represent the relevant transfer of probability with a transfer function, $T : \mathscr{A} \times \Omega \times \Omega \to [0, 1]$, such that $\sum_{\omega \in \Omega} T_E(\omega', \omega) = 1$ for all $\omega' \in \Omega$. For any E and all $\omega, \omega' \in \Omega$, $T_E(\omega', \omega)$ (times 100 percent) specifies the percentage of the total probability mass that is transferred from ω' to ω. It is sometimes assumed—e.g., by Lewis but not by Leitgeb (Leitgeb, 2016)—that $\sum_{\omega \in E} T_E(\omega', \omega) = 1$, so that E bears probability 1 after imaging on it. With T in place, we can formulate the recipe for

[10] A metaphysically deflationary conception of possible worlds has it that a possible world is just a maximally complete set of sentences in some propositional language, instead of a "possible totality of facts."

general imaging. Say that $\boldsymbol{q}$ comes from $\boldsymbol{p}$ by *general imaging* if

$$\boldsymbol{q}(\omega) = \boldsymbol{p}(\omega || E) = \sum_{\omega' \in \Omega} \boldsymbol{p}(\omega') T_E(\omega', \omega)$$

The constraint on T of summing to 1 for each ω' ensures that all probability mass is transferred, so no probability mass is created or destroyed, and the result of imaging is again a pmf. As before, the probability of an event $A \in \mathscr{A}$ can be obtained by summing across $\omega \in A$. Lewis claims that conditionalization and imaging are both minimal revisions, but in different senses. While conditionalization "does not distort the profile of probability ratios, equalities, and inequalities among sentences that imply A," imaging "involves no gratuitous movement of probability from worlds to dissimilar worlds" (1976, p. 142). Lewis proves the following possibility result for (sharp) imaging and the probability of conditionals.

THEOREM 9. *(Lewis, 1976, p. 142) The probability of a Stalnaker conditional with a possible antecedent is the probability of the consequent after imaging on the antecedent:* $\boldsymbol{p}(A > B) = \boldsymbol{p}(B || A) = \boldsymbol{q}(B)$.

More important for the purposes of the present essay is that, as Hannes Leitgeb observes, a result about general imaging due to Peter Gärdenfors can be restated in the language of pooling operators (2016). By *update mechanism*, Leitgeb means a function $U : \mathbb{P} \times \mathscr{A} \to \mathbb{P}$ that maps a probability function and a (non-empty) proposition to a probability function. Gärdenfors shows that general imaging is the unique probabilistic revision method that preserves convex combinations of probability measures. Leitgeb repurposes this result, obtaining the following insight about pooling.

THEOREM 10. *(Cf. Gärdenfors, 1982, Theorem 1) Update by general imaging (with respect to a fixed transfer function T) is the unique update mechanism that commutes with linear pooling with respect to arbitrary coefficients.*

(Here, the transfer function, T, is invariant across priors.) Together with Leitgeb's insight, Gärdenfors' theorem makes showing the next result very easy.

Proposition 12. *Convex IP pooling commutes with general imaging.*

(Commutativity with imaging (CI) could be stated as an axiom for pooling operators.) An analogue of Proposition 10 follows immediately from Proposition 12: convexity is also preserved under general imaging.

Interest in imaging extends beyond natural language semantics and the philosophy of language. Part of the concern with identifying the probability of a conditional with the relevant conditional probability, after all, comes from attempts to give *acceptability* conditions for conditionals. Furthermore, some have argued that, while conditional probability represents *matter-of-fact* supposition in the context of probability, imaging represents *counterfactual* supposition.[11] Lewis himself thought that an interpretation like Stalnaker's is right for *subjunctive* conditionals or counterfactuals, but not for indicative conditionals. Imaging finds crucial employment in James Joyce's account of causal decision theory. Causal relationships are thought to be expressed by subjunctive conditionals, so the probability of such conditionals is of central concern on that view (1999). Baratgin and Politzer contend that empirical evidence indicates that general imaging has some claim as a description of actual revision of probabilistic judgment in dynamic environments (2010). Many authors, however, complain that a philosophically defensible interpretation of the requisite similarity relation among possible worlds has yet to be provided. Some see both the Stalnaker conditional and imaging as questionable shifts from epistemology to metaphysics (Arló-Costa, 2007). Moreover, beginning at least with Ramsey, an alternative line of research attempts to provide acceptability conditions for counterfactuals in terms of belief revision theory, eschewing construals of counterfactuals as bearing truth values (Levi, 1996). Nevertheless, imaging seems to have captured the fancy of many philosophers and others working on conditionals, counterfactual reasoning, and decision theory.

[11] Others, however, have offered more uniform accounts of supposition (e.g., Levi, 1996).

7. Discussion

Propositions 9 and 10 have important implications outside of the context of opinion pooling. While IP models have been widely and convincingly advanced as superior to precise Bayesian representations of uncertainty, standard conditionalization *via* certain learning has been by and large retained as the relevant updating rule (Levi, 1978; Girón and Ríos, 1980). Proposition 10 shows that there is no *mathematical* necessity in that retention for convex Bayesians. For those compelled by Jeffrey's vision of learning, they can have their convex sets of probabilities and their probability kinematics, too.[12] A similar point holds for imaging. Since convexity is preserved under imaging, imaging constitutes a possible "dynamics" for convex Bayesians.[13]

Furthermore, in the precise setting, only linear opinion pooling commutes with imaging. But linear opinion pooling does not commute with Bayesian conditionalization. It follows that no precise pooling method commutes with both imaging and Bayesian conditionalization. In this way, one's hand is forced on the question of updating methods by commitments to pooling methods, and *vice versa*. Not so in the imprecise setting.

Table 1. Summary of Pooling and Updating Commutativity

	Linear	**Geometric**	**Convex IP**
External Bayesianity	X	✓	✓
CJC_W	X	✓	✓
CI	✓	X	✓

Another way to put the point is that, if commutativity of learning and pooling is endorsed *and* more than one updating method is found acceptable (depending on context, say), then there may exist no accommodating precise pooling method.

[12]Though, as Diaconis and Zabell's aforementioned result shows us, in a range of cases there is no mathematical necessity in adopting Jeffrey conditionalization in order to obtain the results of Jeffrey conditionalization.

[13] Though it is not uncontroversial that conditionalization or some other type of updating of represents *learning*. Isaac Levi, for instance, writes, "All conditions of rationality are equilibrium conditions. In a sense they are synchronic conditions [...] Furthermore, in stating conditions of rational equilibrium, no prescription is made regarding the psychological path to be taken in moving from disequilibrium or from one equilibrium position to another. In other words, there are no norms prescribing rational learning processes" (Levi, 1970).

Further limitations issue from results obtained in the literature. For example, suppose commutativity of pooling with both standard Bayesian conditionalization (or Jeffrey conditionalization) and marginalization is endorsed. In the precise setting, we are out of luck. Again, not so in the IP setting, as the results of this chapter in conjunction with those of the previous chapter attest. Philosophical positions that argue for or assume that pooling should be of a particular format are answerable for the limitations of those methods. For instance, in the epistemological debate about peer disagreement, a prominent position encourages peers to "split the difference" between their probabilistic opinions (Elga, 2007). So-called *conciliatory* views on disagreement generally counsel revising opinions in the direction of the dissenting opinions. The revision goes by equal-weight or near equal-weight linear pooling (Christensen, 2009). Some consequences of the failure of commutativity with conditionalization are highlighted in (Russell et al., 2015). As indicated above, Russell, et al. allege that a variant of a diachronic Dutch book can be made against parties following such a policy of disagreement resolution. Similar points can be made regarding other properties of particular pooling methods. For example, neither linear nor geometric pooling preserves probabilistic independence in general (Genest and Wagner, 1987), though convex IP pooling preserves Levi's *confirmational irrelevance*, a generalization of probabilistic independence (Proposition 5). Seidenfeld, Schervish, and Kadane offer a decision-theoretic counterexample to the reasonableness of linear pooling on the basis of its failure to preserve independence (2010). Arguing similarly, Elkin and Wheeler present a variant of a Dutch book argument against resolving disagreements according to the equal weight view (Elkin and Wheeler, 2016). I submit that, not only are IP pooling functions more flexible formal tools, but they admit of stronger normative motivations when various prominent pooling criteria (including the commutativity criteria above) are taken as normative yardsticks.

CHAPTER 4

Another Approach to Consensus and Maximally Informed Opinions with Increasing Evidence

1. Introduction

Merging of opinions results (e.g., Blackwell and Dubins, 1962; Gaifman and Snir, 1982; Schervish and Seidenfeld, 1990; Kalai and Lehrer, 1994; Huttegger, 2015b) have underwritten Bayesian rejoinders to complaints about the subjective nature of personal probability (e.g., Savage, 1954; Schervish and Seidenfeld, 1990; Earman, 1992). Such results establish that sufficiently similar priors achieve consensus in the long run when fed the same increasing stream of evidence. Initial subjectivity, the line goes, is of mere transient significance, giving way to intersubjective agreement eventually. Schervish and Seidenfeld (1990) show that Blackwell and Dubins's classic result (1962) can be extended to certain *sets* of probability functions, while Huttegger (2015b) provides sufficient conditions for merging of opinions for Jeffrey conditioning, a generalization of Bayesian conditionalization. This chapter establishes that Huttegger's merging result for Jeffrey conditioning can in turn be extended to sets of probability functions in analogy to Schervish and Seidenfeld's extension (Section 6). I also extend Huttegger's convergence result (to a "maximally informed opinion") for Jeffrey conditioning to convergence of a *set* of probabilities to a shared maximally informed opinion (Section 7). Convergence to a maximally informed opinion is a (weak) Jeffrey conditioning analogue of Bayesian "convergence to the truth" for conditional probabilities.

A version of this chapter will be published as a paper, coauthored with Michael Nielsen, under the same title.

There are a number of motivations for considering merging and convergence for *sets* of probabilities. For example, according to proponents of *imprecise probabilities* rationality does not always demand numerically precise probability judgments—to say nothing of the descriptive adequacy of such precise judgments (e.g., Levi, 1974; Walley, 1991). In this context, merging and convergence results show that a single agent with a certain sort of imprecise credal state fully expects to uniformly strengthen her point of view towards a consensus among the set (merging) and for the consensus point of view to stabilize (convergence). Another interesting interpretation of such results for sets of probabilities is that they allow us to relax so-called *dynamic coherence*. Instead of demanding updating by (Jeffrey) conditionalization, it is enough that posteriors take any value in some set. As Schervish and Seidenfeld put the point for the Bayesian case, "We do not impose a constraint of (full) dynamic coherence. For consensus, it suffices that the agents use conditional probabilities arbitrarily chosen from a class C enveloped by finitely many (mutually absolutely continuous) distributions. Under the conditions of [Theorem 12], asymptotic certainty follows from static coherence"(1990, p. 402). Or we could consider a social setting, in which sets represent beliefs in some community. In Section 8, I detail another application of this chapter's study to probabilistic opinion pooling with imprecise probabilities (Elkin and Wheeler, 2016; Stewart and Ojea Quintana, 2017b).

2. Preliminaries

We work in the framework of Kalai and Lehrer (1994), for the most part adopting the notation of Huttegger (2015b). The reader should be forewarned that much of the notation is not consistent with previous chapters. There are two reasons for this. The first is that adopting the earlier notation would be unduly cumbersome in the present chapter. The second is that the notation used in this chapter is fairly standard in the context of merging and convergence results in probability theory and facilitates comparison with other, related work. Let Ω be a sample space, a set of elementary events or possible worlds. We let $\mathfrak{F}$ denote a σ-algebra of subsets of Ω, i.e., a set of subsets of Ω

closed under complementation and countable unions. Elements of $\mathfrak{F}$ can be interpreted as events or propositions. For example, Ω could be the set of all infinite sequences of tosses of a coin, and $\mathfrak{F}$ would be the set of relevant coin tossing events. Included in $\mathfrak{F}$ would be propositions describing finite initial segments of a sequence like the first flip landing heads, as well as propositions describing the limiting behavior of the sequence such as $\lim_{n\to\infty} H_n/n = 1/4$, where H_n is the total number of heads on the first n flips. Throughout the chapter uppercase blackboard letters like $\mathbb{P}$ and $\mathbb{Q}$ denote (countably additive) probability measures on $(\Omega, \mathfrak{F})$.[1] An event $A \in \mathfrak{F}$ is said to occur *almost surely*, or a.s., if $\mathbb{P}[A] = 1$.

Let $\mathfrak{E}_1, \mathfrak{E}_2, \ldots$ be an infinite sequence of (finite) partitions of Ω. We suppose that, for any n, $\mathfrak{E}_{n+1}$ is a refinement of $\mathfrak{E}_n$, i.e., every element of $\mathfrak{E}_{n+1}$ is a subset of an element of $\mathfrak{E}_n$. For any n, the partition $\mathfrak{E}_n$ can be thought of as the possible information an agent might receive about the actual world $\omega \in \Omega$. A single flip of a coin determines a binary partition of the set of all infinite sequences of tosses: the coin lands heads on the first toss or the coin lands tails.

Let $\mathfrak{F}_n$ be the algebra generated by $\mathfrak{E}_n$. Besides $\varnothing$ and Ω, $\mathfrak{F}_n$ contains all the unions of elements of $\mathfrak{E}_n$. In the coin tossing example, the algebra $\mathfrak{F}_n$ can be thought of as the set of all coin tossing events up to stage n. For any sequence of coin tosses $\omega \in \Omega$, the agent's information after observing n tosses allows her to determine, for every $E \in \mathfrak{F}_n$, whether $\omega \in E$. The refinement assumption on $\mathfrak{E}_1, \mathfrak{E}_2, \ldots$ implies that $\mathfrak{F}_n \subseteq \mathfrak{F}_{n+1}$, $n = 1, 2, \ldots$. We say that the sequence $\{\mathfrak{F}_n : n \in \mathbb{N}\}$ of σ-algebras is a *filtration*. We assume that this filtration increases to the background σ-algebra $\mathfrak{F}$, i.e., $\sigma\big(\bigcup_{n\geqslant 0} \mathfrak{F}_n\big) = \mathfrak{F}$, so that $\mathfrak{F}$ is generated by a countable sequence of subsets of Ω. The informal idea behind the filtration structure is that information is always increasing and eventually captures all propositions of interest.

For any events $A, E \in \mathfrak{F}_n$ with $\mathbb{P}[E] > 0$, we use the standard ratio definition of conditional probability given an event:

[1] The role of countable additivity in a normative theory of probability judgment is a contentious issue. Its role in all of the results discussed here is not: it is presupposed.

(1) $$\mathbb{P}[A|E] = \frac{\mathbb{P}[A \cap E]}{\mathbb{P}[E]}.$$

For a sub-σ-algebra $\mathfrak{F}_n$ of $\mathfrak{F}$, we will also use the standard definition of conditional probability given $\mathfrak{F}_n$: $\mathbb{P}[A \mid \mathfrak{F}_n]$ is a $\mathfrak{F}_n$-measurable function that satisfies

(2) $$\mathbb{P}[A \cap E] = \int_E \mathbb{P}[A \mid \mathfrak{F}_n](\omega)\mathbb{P}[d\omega]$$

for all $E \in \mathfrak{F}_n$.[2] If we think of ω as the "actual world" regarding coin tosses, then, after observing n tosses, $\mathfrak{F}_n$ is the collection of all propositions that the agent can distinguish as true or false in the actual world. So, just as the conditional probability of A given an event E can be interpreted as the probability of A on the supposition that E is true, so the conditional probability of A given $\mathfrak{F}_n$ can be interpreted as the probability of A in light of the information provided by $\mathfrak{F}_n$.[3] In the present framework, conditional probabilities given sub-σ-algebras can be reduced (almost surely) to conditional probabilities given events; the former merely provide a convenient tool for working around conditioning on null events without adopting extra conventions or assumptions. Since each $\mathfrak{F}_n$ is generated by a finite partition $\mathfrak{E}_n$, for almost every $\omega \in \Omega$ we have

(3) $$\mathbb{P}[A \mid \mathfrak{F}_n](\omega) = P(A \mid E(\omega)),$$

[2] For any sub-σ-algebra $\mathfrak{G}$, the existence of the conditional probability $P[A \mid \mathfrak{G}]$ is guaranteed by the Radon-Nikodym theorem. The uniqueness of $P[A \mid \mathfrak{G}]$ is only almost sure. This means that, in general, there are many measurable functions that satisfy equation (2) and differ from each other on a set of measure 0. To mark this fact, we say that there are different *versions* of $P[A \mid \mathfrak{G}]$. Unlike conditional probabilities given events, there may not exist versions of a conditional probability given a sub-σ-algebra such that $P[A \mid \mathfrak{G}](\omega)$ is a probability measure on $(\Omega, \mathfrak{F})$ for all $\omega \in \Omega$. In this case, we say that $P[A \mid \mathfrak{G}]$ is *irregular*. We discuss regularity more below in connection with the Blackwell-Dubins merging of opinions theorem.

[3] As pointed out by Billingsley (2008, p. 438), this heuristic explanation of conditioning on a sub-σ-algebra breaks down when conditional probabilities are not *proper* (see also Blackwell and Dubins, 1975; Seidenfeld, 2001).

where $E(\omega)$ is the unique cell of $\mathfrak{E}_n$ that contains ω. If $P(E(\omega)) = 0$, then the left-hand side of (3) can be defined arbitrarily.

3. Merging of Opinions

What does it mean to say that two opinions merge, or that $\mathbb{P}$ and $\mathbb{Q}$ agree in the limit? Following Blackwell and Dubins (1962), we adopt the *total variation distance* d as a measure of the distance between $\mathbb{P}$ and $\mathbb{Q}$.

$$d(\mathbb{P}, \mathbb{Q}) = \sup_{A \in \mathfrak{F}} |\mathbb{P}[A] - \mathbb{Q}[A]|$$

So if $\mathbb{P}$ and $\mathbb{Q}$ are within ϵ according to the metric d, then they are within ϵ for every event $A \in \mathfrak{F}$. Now we can formulate a natural sense in which two sequences $\{p_n\}$ and $\{q_n\}$ of probability measures might become (and stay) close. We say that $\{p_n\}$ and $\{q_n\}$ *merge* if

$$d(p_n, q_n) \to 0$$

as $n \to \infty$. If $\{p_n\}$ and $\{q_n\}$ are updates (not necessarily Bayesian) of $\mathbb{P}$ and $\mathbb{Q}$ respectively, we say that $\mathbb{P}$ and $\mathbb{Q}$ merge if $\{p_n\}$ and $\{q_n\}$ do.

The main result of Blackwell and Dubins (1962) concerns merging of conditional probabilities. If learning goes by conditionalization, and the priors $\mathbb{P}$ and $\mathbb{Q}$ do not disagree too drastically, then the Blackwell-Dubins result says $\mathbb{P}$ and $\mathbb{Q}$ must assign probability 1 to merging. To make this precise, we say that $\mathbb{Q}$ is *absolutely continuous* with respect to $\mathbb{P}$, denoted $\mathbb{Q} \ll \mathbb{P}$, if for all $A \in \mathfrak{F}$

$$\mathbb{P}[A] = 0 \implies \mathbb{Q}[A] = 0.$$

So $\mathbb{Q}$ agrees with $\mathbb{P}$ about events bearing probability 0. If $\mathbb{P}$ also agrees with $\mathbb{Q}$ about probability 0 events, so that $\mathbb{P} \ll \mathbb{Q}$, we say that $\mathbb{P}$ and $\mathbb{Q}$ are *equivalent* or *mutually absolutely continuous*. In that case, we have for all $A \in \mathfrak{F}$

$$\mathbb{P}[A] = 0 \iff \mathbb{Q}[A] = 0.$$

It turns out that absolute continuity is sufficient for merging of conditional probabilities.[4] We write $\mathbb{P}_{\mathfrak{F}_n} = \mathbb{P}[\cdot \mid \mathfrak{F}_n]$ and $\mathbb{Q}_{\mathfrak{F}_n} = \mathbb{Q}[\cdot \mid \mathfrak{F}_n]$.[5]

THEOREM 11. *(Blackwell and Dubins, 1962) If $\mathbb{Q} \ll \mathbb{P}$, then $d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \to 0$ a.s. $(\mathbb{Q})$ as $n \to \infty$.*

With $\mathbb{Q}$-probability 1, $\mathbb{P}$ merges to $\mathbb{Q}$ when $\mathbb{Q} \ll \mathbb{P}$. If $\mathbb{P} \ll \mathbb{Q}$ also, we have merging from the perspectives of both $\mathbb{P}$ and $\mathbb{Q}$. The almost surely qualifications are needed in the statement of Theorem 11 because the conditional probabilities $\mathbb{P}_{\mathfrak{F}_n}$ and $\mathbb{Q}_{\mathfrak{F}_n}$ in this general setting are random objects, that is, they depend on $\omega \in \Omega$ (see footnote 4 for details).

4. Merging of Opinions for Sets of Probabilities

Schervish and Seidenfeld provide sufficient conditions to secure merging of opinions for *sets* of probability functions. A set of probability functions C is called *convex* if $\mathbb{P}, \mathbb{Q} \in C$ implies $\alpha\mathbb{P} + (1-\alpha)\mathbb{Q} \in C$ where $\alpha \in [0,1]$. That is, for any two elements of C, C includes all convex combinations of those elements. We call $\mathbb{P} \in C$ an *extreme point* of C if $\mathbb{P} = \alpha\mathbb{Q} + (1-\alpha)\mathbb{R}$ with $\mathbb{Q}, \mathbb{R} \in C$ and $\alpha \in [0,1]$ implies $\mathbb{Q} = \mathbb{P}$ or $\mathbb{R} = \mathbb{P}$. The following result is a consequence of Theorem 11.

THEOREM 12. *(Schervish and Seidenfeld, 1990, Corollary 1) Let C be a closed, convex set of probability functions all mutually absolutely continuous, and generated by finitely many extreme*

[4]Kalai and Lehrer show that absolute continuity is not merely a sufficient condition for merging of conditional probabilities, it is also a necessary one (1994, Theorem 2).

[5]Besides absolute continuity, the Blackwell-Dubins theorem also requires that the conditional probabilities $\mathbb{P}_{\mathfrak{F}_n}$ and $\mathbb{Q}_{\mathfrak{F}_n}$ be *regular* (or what Blackwell and Dubins call *predictive*): $\mathbb{P}_{\mathfrak{F}_n}(\omega)$ and $\mathbb{Q}_{\mathfrak{F}_n}(\omega)$ are probability measures on $(\Omega, \mathfrak{F})$ for each $\omega \in \Omega$. It follows from the fact that each $\mathfrak{F}_n$ is generated by a finite partition that $\mathbb{P}_{\mathfrak{F}_n}$ and $\mathbb{Q}_{\mathfrak{F}_n}$ have regular versions. We assume throughout that all conditional probabilities are regular. For examples of conditional probabilities that are irregular, see Billingsley (2008, Exercise 33.11) and Seidenfeld (2001).

points. Then, almost surely, the conditional probabilities of elements of C merge uniformly, that is, $\sup_{\mathbb{P},\mathbb{Q}\in C} d(\mathbb{P}_{\mathfrak{F}_n}(\omega), \mathbb{Q}_{\mathfrak{F}_n}(\omega)) \to 0$ *as* $n \to \infty$ *for almost all* $\omega \in \Omega$.

Another way of stating the conclusion of Theorem 12 is that for almost all ω and all $\epsilon > 0$ there is some m such that for all $n \geqslant m$ and all $\mathbb{P}, \mathbb{Q} \in C$, $d(\mathbb{P}_{\mathfrak{F}_n}(\omega), \mathbb{Q}_{\mathfrak{F}_n}(\omega)) < \epsilon$. That the same m works for *all* $\mathbb{P}$ and $\mathbb{Q}$ is what it means for merging to be *uniform*.

Here is the gist of their proof. Since C is generated by finitely many extreme points, the (total variation) distance between the conditional probabilities of any of the extreme points of C is bounded by the maximum distance for pairs of conditional probabilities of extreme points. Moreover, that maximum is a bound for the distance between the conditional probabilities of any points in the whole set. To see this, note two things. First, $d(\mathbb{P}, \mathbb{Q})$ bounds the distances $d(\mathbb{P}, \mathbb{R})$ and $d(\mathbb{R}, \mathbb{Q})$ where $\mathbb{R}$ is a convex combination of $\mathbb{P}$ and $\mathbb{Q}$. Second, at any stage n, the conditional probabilities of points in C can be written as convex combinations of the conditional probabilities of the extreme points, i.e., there exist $\alpha_1, ..., \alpha_k$ with $\alpha_i \in [0, 1]$ and $\sum_{i=1}^{k} \alpha_i = 1$ such that $\mathbb{R}_{\mathfrak{F}_n} = \sum_{i=1}^{k} \alpha_i \mathbb{P}^i_{\mathfrak{F}_n}$ for all $\mathbb{R} \in C$ and all n. Theorem 12 is stronger than it may at first appear. Notice that *any* set of probabilities that is included in a closed, convex set generated by finitely many extreme points merges uniformly by this result.

It is not easy to relax the assumption in Theorem 12 that C is convex and generated by finitely many extreme points. Schervish and Seidenfeld demonstrate this with several other results and examples. For instance, if C is d-compact, an assumption that is weaker than the assumptions of Theorem 12, the conditional probabilities of elements of C may not exhibit almost sure uniform merging. However, Schervish and Seidenfeld show that weaker modes of merging can be achieved under weaker assumptions about C, such as d-compactness (see their Corollaries 2 and 3). It is worth noting that any negative result for Bayesian conditionalization and merging carries over to Jeffrey conditioning, a generalization of conditionalization introduced and studied below. An immediate consequence of Schervish and Seidenfeld's results, then, is that Jeffrey updates of a d-compact set of probability measures may not almost surely merge uniformly.

There are a number of uses to which results such as Theorem 12 might be put. We delay discussing some of the significance of Theorem 12 and Propositions 13 and 14 (to come) until Section 8.

5. Jeffrey Conditioning

Bayesian conditionalization is a putative diachronic norm that specifies how probabilistic learning takes place.[6] A basic assumption of conditionalization is that it is *propositions* (or *events*) that are learned. In other words, learning experiences can always be represented by some $E \in \mathfrak{F}$. When an agent learns E, conditionalization says that her *posterior* probability $\mathbb{P}_E[A]$ for an event $A \in \mathfrak{F}$ should be equal to her prior conditional probability of A given E (provided, of course, that this is well defined), i.e.

$$\mathbb{P}_E[A] = P[A \mid E].$$

It is clear from (1) that conditionalization requires that E bear posterior probability 1. Jeffrey conditioning relaxes both of the fundamental features of conditionalization: it does not assume that learning experiences are always represented by propositions, and it does not require that learning involves assigning propositions probability 1.

Proponents of Jeffrey conditioning want to allow for "uncertain learning." Uncertain learning induces a change in the probabilities P_n assigned to the members of the partition $\mathfrak{E}_n$. Jeffrey conditioning applies just in case the change over $\mathfrak{E}_n$ is *rigid*, i.e.

[6] In this essay, I will not quarrel with the view that conditionalization and Jeffrey conditioning are genuine diachronic norms (outside of indicating how they can be relaxed while still achieving merging and convergence). However, I regard diachronic "learning" norms with a good deal of suspicion. I recognize that this is likely at odds with most philosophical writing on probability. The purely synchronic construal of merging views the phenomenon as one related to supposition.

$$\mathbb{P}_n[A|E] = \mathbb{P}_{n-1}[A|E], \text{ for all } E \in \mathfrak{E}_n \text{ and all } A \in \mathfrak{F}. \tag{4}$$

The rigidity condition (4) says that the update from $\mathbb{P}_{n-1}$ to $\mathbb{P}_n$ does not change conditional probabilities given members of the partition $\mathfrak{E}_n$. This is equivalent to the requirement that for all members E of the partition $\mathfrak{E}_n$ and all subsets A, B of E, the update from $\mathbb{P}_{n-1}$ to $\mathbb{P}_n$ does not change the ratio of the probabilities of A and B. The law of total probability and (4) yield the familiar Jeffrey conditioning equation, which extends P_n from $\mathfrak{E}_n$ to the entire algebra $\mathfrak{F}$:

$$\mathbb{P}_n[A] = \sum_{E \in \mathfrak{E}_n} \mathbb{P}_{n-1}[A|E] P_n[E], \text{ for all } E \in \mathfrak{E}_n \text{ and all } A \in \mathfrak{F}. \tag{5}$$

If $\mathfrak{E}_n = \{E, E^c\}$ and $P_n[E] = 1$, then Jeffrey conditioning reduces to standard conditionalization, i.e. $\mathbb{P}_n[A] = \mathbb{P}_{n-1}[A \mid E]$ for all $A \in \mathfrak{F}$.[7] Equation (5) simplifies in another way under our assumption that the sequence of partitions $\mathfrak{E}_1, ..., \mathfrak{E}_n$ is such that $\mathfrak{E}_{i+1}$ refines $\mathfrak{E}_i$ for all $1 \leqslant i \leqslant n-1$. In this case, after $n-1$ applications of equation (4) in equation (5), we have

$$\mathbb{P}_n[A] = \sum_{E \in \mathfrak{E}_n} \mathbb{P}[A|E] P_n[E], \text{ for all } E \in \mathfrak{E}_n \text{ and all } A \in \mathfrak{F}. \tag{6}$$

Equation (6) shows that, in our framework, posterior probabilities $\mathbb{P}_n$ are determined by their values on members of the partition $\mathfrak{E}_n$ and prior ($\mathbb{P}$) conditional probabilities given members of $\mathfrak{E}_n$.

[7] This, of course, requires that we adopt some conventions for conditioning on null events, otherwise the left-hand side of equation (4) may be undefined.

Unless otherwise stated, we will be assuming that the events $E \in \mathfrak{E}_n$ have positive prior probability so that the right-hand side of (6) is well defined.

It is worth pointing out that, in certain situations, Jeffrey conditioning can be represented as standard Bayesian conditionalization in a richer algebra. The cases in which such so-called *superconditioning* is possible are characterized by Diaconis and Zabell's superconditioning criterion (1982, Theorem 2.1). But the superconditioning criterion is not trivial; it fails to hold in many cases. Hence, Huttegger's merging result for Jeffrey conditioning in the following section and our merging result for sets of Jeffrey updates in Section 6 genuinely generalize Blackwell and Dubins's and Schervish and Seidenfeld's results, respectively. (For other reservations about reducing Jeffrey conditioning to Bayesian conditionalization in the setting of merging of opinions, see Huttegger's discussion (2015b, pp. 630–631).)

5.1. Merging. Huttegger proves an analogue of Blackwell and Dubins's merging result for Jeffrey conditioning. Just as we considered random conditional probabilities in the case of conditionalization, we now treat posteriors over the learning partition as random. By random probability, we mean $P_n[\cdot]$ is not a determinate number but rather a (measurable) function on Ω such that $P_n[\cdot](\omega)$ is a probability measure on $\mathfrak{E}_n$ for each $\omega \in \Omega$. Let $p_n^E = P_n[E] = \mathbb{P}_n[E]$ be a random probability assigned to $E \in \mathfrak{E}_n$. We treat future probabilities as random quantities because, from an agent's present point of view concerning $E \in \mathfrak{E}_n$, the precise value p_n^E is unknown.

Because Jeffrey conditioning is less constrained than Bayesian conditionalization, some additional assumptions figure into Huttegger's theorem. The first additional assumption is that the sequence $\{Q_n\}$ is *uniformly absolutely continuous* with respect to $\mathbb{Q}$. This condition consists of two requirements: (i) $Q_n \ll \mathbb{Q}|_{\mathfrak{F}_n}$ for all n, and (ii) for every $\epsilon > 0$ there is a $\delta > 0$ such that for all n and all $A \in \mathfrak{F}_\mathfrak{n}$,

$$\mathbb{Q}[A] < \delta \implies Q_n[A] < \epsilon.$$

(i) is the usual constraint on Bayesian learning that events that are null for the prior $\mathbb{Q}$ remain null for the posterior $\mathbb{Q}_n$. The additional uniformity requirement (ii) ensures that this relation holds as we pass to the limit (Huttegger, 2015b, p. 623).[8]

The other assumption demands a certain sort of stability of probability judgments as the agent updates over increasingly fine partitions. While we treat posterior probabilities as random, condition (M') places significant constraints on sequences of probability judgments. For all n and all $E \in \mathfrak{E}_n$,

$$\text{(M')} \qquad \int_G p^E_{m+1} d\mathbb{P} = \int_G p^E_m d\mathbb{P} \text{ for all } m \geqslant n \text{ and all } G \in \mathfrak{F}_m$$

Property (M') is the *martingale* condition, and requires that future probabilities are equal to present probabilities on average. For all n and all $E \in \mathfrak{E}_n$, the sequence $p^E_n, p^E_{n+1}, \ldots$ must form a martingale. The martingale condition has been defended as an essential feature of rational learning experiences (Huttegger, 2013, 2015b). In the case of Bayesian conditionalization, future probabilities are fixed at 1 for those events learned. This, of course, is not generally the case for Jeffrey conditioning.

Theorem 13 is Huttegger's merging result for Jeffrey conditioning.

THEOREM 13. *(Huttegger, 2015b, Theorem 9.2) Suppose that* $P_n, n = 1, 2, \ldots$ *and* $Q_n, n = 1, 2, \ldots$ *are random sequences of probability measures on* $(\Omega, \mathfrak{F}_1), (\Omega, \mathfrak{F}_2), \ldots$ *with* $Q_n = P_n$ *a.s., that* $Q_n, n = 1, 2, \ldots$ *is uniformly absolutely continuous with respect to* $\mathbb{Q}$ *a.s. ($\mathbb{Q}$), and that* $\mathbb{Q} \ll \mathbb{P}$*. If (M') holds for* $Q_n, n = 1, 2, \ldots$ *a.s. ($\mathbb{Q}$), then* $d(\mathbb{P}_n, \mathbb{Q}_n) \to 0$ *as* $n \to \infty$ *a.s. ($\mathbb{Q}$).*

[8] Here is another way of thinking about the uniform absolute continuity requirement. First, if $Q_n \ll \mathbb{Q}|_{\mathfrak{F}_n}$ for all n, then $\{Q_n\}$ is uniformly absolutely continuous with respect to $\mathbb{Q}$ if and only if the sequence $\{Y_n\} = \{dQ_n/d\mathbb{Q}\}$ of Radon-Nikodym derivatives is uniformly integrable. The "only if" direction of this statement is shown by Huttegger (2015b, Lemma 12.1), and the "if" direction is not difficult to prove. The uniform integrability of $\{Y_n\}$ in turn implies that $\{Y_n\}$ is uniformly bounded in expectation, i.e. $\sup_n \mathbb{E}[Y_n] < \infty$ (again, this is a standard result). Therefore, if we think of the derivatives Y_n as representing the "rate" at which the posteriors Q_n are changing with respect to the prior $\mathbb{Q}$, then Huttegger's uniform absolute continuity condition demands that this rate is not expected to diverge to ∞.

6. Merging of Opinions for Jeffrey Conditioning and Sets of Probabilities

We are now in a position to turn to the contributions of the present chapter. The first result extends Huttegger's theorem for merging of opinions *via* Jeffrey conditioning to sets of probabilities, an analogue of Theorem 12.

Proposition 13. *Let C be a closed, convex set of probability functions all mutually absolutely continuous, and generated by finitely many extreme points. Suppose that almost surely each pair of probabilities $\mathbb{P}, \mathbb{Q} \in C$ satisfies the conditions of Theorem 13. Then almost surely the elements of C merge uniformly.*

A proof is included in the appendix. As Schervish and Seidenfeld indicate, Theorem 12 is a simple corollary of Theorem 11 given that point-by-point conditionalization preserves the convexity of a set. Proposition 13 is especially interesting in light of the fact that, unlike Bayesian conditionalization, Jeffrey conditioning does not generally preserve the convexity of the initial set—even for convex sets of priors generated by finitely many extreme points. I provide a demonstration in the appendix (Proposition 27).

7. Maximally Informed Opinions

Another consequence of Doob's martingale convergence theorem is the so-called "convergence to truth" or "convergence to certainty" for conditional probabilities, captured by the following statement:

$$\text{For all } A \in \mathfrak{F},\ \lim_{n\to\infty} \mathbb{P}_{\mathfrak{F}_n}[A] = \mathbf{1}_A(\omega)\ a.s.\ (\mathbb{P}).$$

Put another way, coherence requires assigning probability 1 to converging to "the truth" for any event $A \in \mathfrak{F}$ whose truth value is determined by observations. One way of interpreting such convergence is that "[i]t shows that evidence triumphs over prior opinions under the appropriate circumstances" (Huttegger, 2015a, p. 590). Two mutually absolutely continuous priors assign probability 1 not just to merging, but to being certain of the same observational hypotheses in the limit.[9]

Convergence to certainty does not obtain in general for Jeffrey conditioning. However, Huttegger shows that Jeffrey conditioning *does* lead probabilities to settle down to *some* limit if we assume (M'), just not necessarily to 0 or 1.[10] Skyrms labels this limit a "maximally informed opinion" (1996).

Convergence to maximally informed opinions can be extended to sets of probabilities. All probabilities in C converge to the same limit, $\mathbb{P}_\infty$, almost surely. To see this, consider that all elements in C almost surely are converging to a limit (Huttegger, 2015b, Theorem 9.1) *and* are merging (Theorem 13). If two elements $\mathbb{P}, \mathbb{Q} \in C$ converged to *different* limits (with positive probability), then there would be some ϵ and some m such that $d(\mathbb{P}_n, \mathbb{Q}_n) > \epsilon$ for any $n \geqslant m$, which is inconsistent with merging. In fact, a slightly stronger observation can be demonstrated using our Proposition 13. As indicated in the introduction, dynamic coherence can be relaxed. For convergence to a maximally informed opinion, it suffices that probabilities take their values in the set of Jeffrey updates of C.

Proposition 14. *Let C be as in Proposition 13. If $f_n \in \{\mathbb{P}_n : \mathbb{P} \in C\}$, then the sequence f_n converges to the maximally informed opinion $\mathbb{P}_\infty$ a.s. ($\mathbb{P} \in C$).*

We omit the proof.

8. An Application: Pooling and Learning, and Learning, and Learning ...

Earlier, I indicated some consequences of these merging results for sets of probabilities. First, if an agent has a credal state properly represented by a set of probability functions, Theorem

[9]While Cisewski et al., for example, refer to such convergence results as "desirable" and "laudable" (MS), others are less enthusiastic (e.g., Earman, 1992; Kelly, 1996; Belot, 2013).

[10]For a discussion of this point, see (Huttegger, 2015b, §7). In particular, Huttegger's Theorem 7.1 (and its generalization, Theorem 9.1) states that sequences of uniformly absolutely continuous Jeffrey updates that satisfy (M') converge setwise, i.e. for all $A \in \mathfrak{F}$, $\lim_{n\to\infty} \mathbb{P}_n[A]$ exists (almost surely, in the case of Theorem 9.1). In fact, something stronger holds: the set function $\mathbb{P}_\infty$ defined by $\mathbb{P}_\infty[A] = \lim_{n\to\infty} \mathbb{P}_n[A]$ is a probability measure. This follows from the Vitali-Hahn-Saks theorem. See Nielsen (MS) for further discussion.

12 and Proposition 13 provide sufficient conditions for that indeterminacy to be reduced by evidence in the limit, for both Bayesian conditionalization and Jeffrey conditioning. Second, Theorem 12 and Proposition 13 have interesting implications for issues related to dynamic coherence as Schervish and Seidenfeld point out. To achieve asymptotic consensus, agents need not be dynamically coherent—they need not conditionalize or update by Jeffrey conditioning. Interpreting merging results as claims about agents' unconditional probabilities through time, it is sufficient to achieve asymptotic consensus that posteriors take their values in some appropriate set (the set that results from updating all elements of C). A similar point holds for relaxing dynamic coherence for convergence.

A third upshot reinterprets the first point above in a social setting. The probabilistic aggregation framework studied in previous chapters is a general and precise setting in which to study ways of forming a *consensus* or group point of view from a set of potentially diverse points of view. Different concrete recipes for pooling probability judgments have been studied.[11] Some of the most extensively studied are ways of *averaging* probabilities to arrive at a group probability. For example, we could take linear or geometric averages of profiles of individual probabilities. In general, pooling allows us to consider ways of arriving at a "consensus" besides those of updating on a shared infinite stream of evidence. One assumption common throughout the pooling literature is that pooling produces a single, "group" or "consensus" probability function. Yet the standard frameworks have significant limitations. A number of results show that certain sets of desirable aggregation properties cannot be simultaneously satisfied on pain of triviality or inconsistency. Drawing on work on imprecise probabilities, Chapter 2 motivates the use of imprecise probabilities in the context of pooling and generalize the canonical mathematical framework to allow for set-valued pooling functions. A number of simple possibility results were established and a distinguished format of pooling with imprecise probabilities is characterized.[12]

[11] For a good survey, see (Genest and Zidek, 1986).

[12] The framework of the preceding chapters is general, however, subsuming any mapping from a profile of probabilities to a set of probabilities.

That format is the function that returns the convex hull of any profile of probabilities to be aggregated. The point of departure for the concern with the convex hull—and the use of imprecise probabilities in the context of pooling more generally—is an essay of Isaac Levi's (1985). There, Levi distinguishes between *consensus as shared agreement*, which is available at the outset of inquiry by retaining agreements and suspending judgment on other matters, and *consensus as the outcome of inquiry*, which emerges when evidence resolves initial disputes. About the role of the convex hull, Levi writes,

> a potential resolution of the conflict between rival credal probability distributions is to be represented by a credal probability distribution which is the weighted average of the distributions in conflict. Hence, the set of all potential resolutions of such a conflict is to be represented by the convex hull (the set of all weighted averages) of the credal distributions initially in contention. My assumption was that this convex set of probability distributions represented the first kind of consensus I regard as important—consensus as shared agreements regarding probability judgment (Levi, 1985, p. 6).

On this view, a pooling function that returns the convex hull can be regarded as delivering a consensus position at the outset of inquiry.

One potential complaint about such a view of pooling is that, in many cases, it would identify rather *weak* consensus positions. According to Levi, this is as consensus at the outset of inquiry should be. But that is not the end of his story. Weak points of view can be strengthened through inquiry. Levi writes,

> if we adopt a consensus as shared agreement on credal probabilities before the acquisition of evidence, we may hope to obtain new data via experimentation and observation which will yield a consensus which resolves the original dispute via inquiry. In typical cases, *ample data will lead via Bayes's theorem and conditionalization to a reduction in the indeterminacy in the state of credal probability judgment. Consensus as the outcome of inquiry will be more determinate than the consensus as shared agreements adopted at the outset of inquiry* (Levi, 1985, p.10, emphasis ours).

One might see merging results as validating Levi's claim here. Under certain conditions, conditionalization does lead to a reduction in indeterminacy. And, happily enough, those conditions are met when forming the convex hull of some finite profile of probability functions to be aggregated into an initial consensus position. Furthermore, our Proposition 13 shows that the ability to achieve the desired reduction of indeterminacy is not limited to Bayesian conditionalization. Under a few further assumptions, Jeffrey conditioning too reduces indeterminacy uniformly as evidence accumulates.[13] Levi's picture of consensus appears to be somewhat robust against the choice of updating method. It is also worth emphasizing again that any pooling function that outputs a subset of the convex hull of the profile presents a consensus position that is subject to all of the merging and convergence results for sets mentioned in this essay. If Levi is right that the convex hull represents "the set of all potential resolutions" of the conflict in probability judgments, and reasonable pooling functions take values in the power set of the set of such resolutions, then the merging and convergence results discussed here hold for the class of reasonable pooling functions. Merging results, I submit, constitute a partial response to complaints concerning the use of imprecise probabilities to identify a consensus at the outset of inquiry in the context of pooling.

[13] With very few exceptions (e.g., Stewart and Ojea Quintana, 2017a), "uncertain learning" is not treated in the context of uncertainty that is not reducible to a numerically precise probability.

CHAPTER 5

Unanimous Consensus against AGM?

1. Introduction

Concern for epistemic consensus is a hallmark of pragmatist epistemology. In "The Fixation of Belief," Peirce criticizes certain methods of belief formation on the grounds that their failure to secure wide-spread consensus will undermine the resulting beliefs (Peirce, 1992a, pp. 116-117). In other places, Peirce can be interpreted as proposing a definition of truth in terms of consensus in the long run of inquiry, or as asserting that, if inquiry were to go on indefinitely, the truth would be consensually settled upon in the limit (e.g., Peirce, 1992b, pp. 138-139; Misak, 2004, pp. 67-70). Later, and notoriously, Rorty advocated consensus-type accounts of knowledge and truth. As Guignon and Hartely summarize his view, "There is no basis for deciding what counts as knowledge and truth other than what one's peers will let one get away with in open exchange of claims, counterclaims, and reasons" (2003, p. 11). Consensus also figures prominently in Isaac Levi's brand of pragmatism. Levi denies that "inquiry can proceed without appeal to some point of view (state of full belief, demands for information, judgments of credal probability, etc.)," while also denying that "there is some standard, objective point of view to which appeal may always be made" (1991, p. 87). However, these denials are not tantamount to "cognitive licentiousness" or "epistemic anarchy," according to Levi, because we can always identify a *consensus as shared agreement* that can function as a non-question-begging initial position for joint inquiry (1991, pp. 87-88). On Levi's proposal, parties to a consensus should restrict themselves to the shared agreements between their points of view. Once on common ground, the group can engage in hypothetical reasoning and inquiry.

Unanimity or *shared agreement* is one simple and perhaps obvious account of consensus. On this view, epistemic consensus among some belief sets consists of the beliefs held in common, those beliefs that are unanimously held (Levi, 1996, Chp. 2). In a sense, shared agreement is the analogue of the conciliatory position in the literature on peer disagreement for sets of beliefs (e.g., Christensen, 2009; Feldman, 2011). A shift to consensus as shared agreement suspends judgment on beliefs about which there is not agreement. It turns out, though, that on the received view of how an agent ought to revise her beliefs, AGM belief revision theory, unanimous consensus is *not* always available to serve as the basis of joint inquiry or deliberation. The key move here is to think of reaching a consensus (for an agent) as an agent revising her beliefs to a belief set that represents a consensus position for some set of rival belief sets. In light of the results presented below, we are confronted with a classic *modus ponens/modus tollens* dilemma. On the one hand, if an agent should always be able to contract her belief set to unanimous consensus with another belief set, then the propositions that follow are damning for AGM partial meet contraction. On the other hand, if the AGM account or another account for which analogous results hold is found sufficiently compelling, so much the worse for the availability of unanimity for the purposes of collective inquiry and deliberation. This note establishes the existence of such a dilemma.

2. AGM Contraction

The unanimous consensus or consensus as shared agreement for two sets of beliefs, K_1 and K_2, is given by the beliefs common to both sets, that is, by the intersection: $K_1 \cap K_2$. So, a transition from either K_1 or K_2 to $K_1 \cap K_2$ must be a *contraction*, an operation of belief removal. How ought an agent to go about contracting her beliefs? The AGM paradigm offers perhaps the most familiar and well-explored proposal. Let's briefly review it.

Let $\mathcal{L}$ be a propositional language that is closed under truth functional operations. We let lower case Greek letters, $\alpha, \beta, ...$ (except γ), range over sentences, and capital Roman letters, $A, B, ...$, refer to *sets* of sentences. By closure, if $\alpha \in \mathcal{L}$, then $\neg\alpha \in \mathcal{L}$. If $\alpha, \beta \in \mathcal{L}$, then $\alpha \vee \beta$ is in $\mathcal{L}$, and so on. We use lower case Roman letters, $p, q, ...$, to denote the atomic formulae of the language. Let Cn be a Tarskian consequence operator: $Cn : \mathscr{P}(\mathcal{L}) \to \mathscr{P}(\mathcal{L})$. A sentence α is in $Cn(A)$ if and only if α is a logical consequence of A.[1] We call K a *theory* if K is deductively closed, i.e., $K = Cn(K)$. Let $\mathbb{K}$ denote the set of all deductively closed subsets of $\mathcal{L}$. Belief sets are elements of $\mathbb{K}$, with sentences representing beliefs.

There are three standard belief change operations. In expansion, a sentence is added to K. But expansion can introduce inconsistency. Revision incorporates a sentence into K while preserving consistency. Contraction removes beliefs from K. We put $K \dot{-} \alpha$ for the result of contracting α from K with $\dot{-} : \mathbb{K} \times \mathcal{L} \to \mathbb{K}$. There are choices to be made in contraction since, in general, how to alter K so that a sentence is no longer a consequence is not uniquely determined. For example, if $\alpha, \alpha \to \beta, \beta \in K$, then at least one of the other sentences must also be surrendered if β is to be contracted. The AGM postulates for contraction are as follows.

($\dot{-}$1) $K \dot{-} \alpha = Cn(K \dot{-} \alpha)$ whenever $K = Cn(K)$

($\dot{-}$2) $K \dot{-} \alpha \subseteq K$

($\dot{-}$3) If $\alpha \notin K$, then $K \subseteq K \dot{-} \alpha$

($\dot{-}$4) If $\alpha \in K \dot{-} \alpha$, then $\alpha \in Cn(\varnothing)$

($\dot{-}$5) $K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$

($\dot{-}$6) If $Cn(\alpha) = Cn(\beta)$, then $K \dot{-} \alpha = K \dot{-} \beta$

($\dot{-}$7) $K \dot{-} \alpha \cap K \dot{-} \beta \subseteq K \dot{-} (\alpha \wedge \beta)$

($\dot{-}$8) If $\alpha \notin K \dot{-} (\alpha \wedge \beta)$, then $K \dot{-} (\alpha \wedge \beta) \subseteq K \dot{-} \alpha$

The first six postulates are known as the *basic* postulates; the last two, as the *supplementary* postulates.

[1]As is customary, we assume Cn satisfy certain standard properties:

(Inclusion)	$A \subseteq Cn(A)$
(Monotonicity)	If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$
(Iteration)	$Cn(A) = Cn(Cn(A))$
(Supraclassicality)	If α is a classical consequence of A, then $\alpha \in Cn(A)$
(Deduction)	$\beta \in Cn(A \cup \{\alpha\})$ *iff* $\alpha \to \beta \in Cn(A)$
(Compactness)	If $\alpha \in Cn(A)$, then $\alpha \in Cn(A')$ for some finite $A' \subseteq A$

Various concrete contraction constructions have been explored in the literature. In *partial meet* contraction, the central focus is on so-called *remainder sets*, the set of all maximal subsets of K that do not imply α, the α-*remainders* of K. The typical stated motivation for such focus is informational economy. Because information is valuable, agents should seek to retain as much of it as possible when contracting their belief sets.[2]

Definition 1. The set of α-remainders of K is given by

$$K \perp \alpha = \{K' \subseteq K : (i)\ \alpha \notin Cn(K'), (ii) \text{ if } \beta \in K \text{ but } \beta \notin K',\ \text{then } \beta \rightarrow \alpha \in K'\}$$

A partial meet contraction takes the intersection of some selection of elements in $K \perp \alpha$. The selection function $\gamma : \mathscr{P}(\mathbb{K}) \rightarrow \mathscr{P}(\mathbb{K})$ chooses a nonempty subset $\gamma(K \perp \alpha)$ of $K \perp \alpha$ (when $K \perp \alpha = \varnothing, \gamma(K \perp \alpha) = K$). We call $\dot{-}$ a partial meet contraction operator just in case there exists a selection function such that the following holds:

$$K \dot{-} \alpha = \bigcap \gamma(K \perp \alpha)$$

The six basic contraction postulates characterize partial meet contraction.

Theorem 14. *(e.g., Gärdenfors, 1988) A contraction function, $\dot{-}$, is a partial meet contraction function iff $\dot{-}$ satisfies* $(\dot{-}1) - (\dot{-}6)$.

The supplementary postulates constrain the partial meet contractions further (by constraining the behavior of γ).

[2] Levi has stressed in numerous places that the loss of informational value—and not of informational content—is what should be minimized in contraction (1991; 2004). The point here is that it is possible for K' to bear as much informational value as K even if $K' \subset K$. Consequently, restricting our attention to remainder sets is misguided. Hans Rott argues that, in any case, informational economy does not play the role in AGM belief revision that it is generally ascribed (2000). Consider, for example, the fact that partial meet contraction and not maxichoice contraction plays the central role.

3. Contracting to Consensus?

Articulating an account of how to transition to a desired corpus has not been the AGM agenda. Instead, the aim is to provide an account of how to accommodate some given input (in contraction, the "input" sentence is to be removed). The motivation or justification for the input sentence is not part of the formal account of belief revision in the AGM paradigm. But one motivation for contracting a particular sentence may be to achieve consensus with another corpus. That is, the desire to shift to a specific sub-corpus may underwrite contracting a particular sentence. When two agents with consistent belief sets, K_1 and K_2, disagree on some δ in the sense that $\delta \in K_1$ and $\neg\delta \in K_2$, it is clear that each can suspend judgment *on* δ by contracting by δ and $\neg\delta$, respectively (Elkin, 2015). Is focusing on a single sentence in general sufficient to reach consensus as shared agreement? Put another way, is the belief set that represents unanimous consensus with another belief set always accessible in the AGM framework?

No. We can establish the following proposition.

PROPOSITION 15. *It is* ***not*** *the case that, for any consistent belief sets* $K_1, K_2 \in \mathbb{K}$ *and any partial meet contraction operator* $\dot{-}$*, there is a sentence* α *such that* $K_1 \dot{-} \alpha = K_1 \cap K_2$*.*

(Proofs of propositions are relegated to the appendix as they are throughout the dissertation .) The upshot of Proposition 15 is that there are partial meet contraction operators for which consensus with certain other belief sets is *not* available. We might wonder, however, about epistemically opportunistic agents. Perhaps an agent would be willing to employ a *different* partial meet contraction operator just to reach consensus. But much of the work in belief revision presupposes that an agent is committed to a unique contraction operator at a given time (e.g., Hansson, 2003, p. 44). Given well-known results in the literature, this is tantamount to an agent's being committed (at a given time) to a single way of assessing what the "best" elements of $K\perp\alpha$ are in the case of partial meet contraction, or a single way of assessing the epistemic usefulness or value of sentences in $\mathcal{L}$ in the case of epistemic entrenchment, or a single way of assessing, essentially, the plausibility of possible worlds in the case of a Grove system of spheres.

Furthermore, in the infinite case, there are other straightforward limitations, as the next proposition attests. Proposition 16 demonstrates that the strategy of either putting further restrictions on the class of partial meet contractions—so that only full meet or some other subset of the partial meet contractions are allowed—or permitting agents to employ alternative partial meet contraction operators for the purpose of contracting to consensus is not a general workaround for the limitation indicated in Proposition 15.

PROPOSITION 16. *Let $\mathcal{L}$ contain infinitely many atomic sentences. It is **not** the case that for every pair of consistent belief sets $K_1, K_2 \in \mathbb{K}$ there is some partial meet contraction, $\dot{-}$, such that $K_1 \dot{-} \alpha = K_1 \cap K_2$ for some sentence α.*

So there are belief sets for which *no partial meet contraction operator* yields consensus for some sentential input. Proposition 16 represents more than a mere remote mathematical obstacle for AGM contraction according to a pragmatist of Levi's stripe. Pragmatists have long resisted accounts of privileged or fixed languages and conceptual schemes. We should be open to refining our language as appropriate.

We might consider equipping each belief set with a special-purpose consensus contraction operator, $\dot{-}_{C_{K_2}}$, such that $K_1 \dot{-}_{C_{K_2}} \alpha = K_1 \cap K_2$ for any $\alpha \in \mathcal{L}$. In a similar vein, Levi considers *consensus-based revisions* which result from first contracting to shared agreements (e.g., $K_1 \dot{-}_{C_{K_2}} \alpha$), then adding a sentence (Levi, 1996, p. 42). Consensus-based revisions can be used to engage in group hypothetical reasoning from a shared background corpus, for example. As the next observation shows, $\dot{-}_{C_{K_2}}$ is at odds with the AGM vision of belief contraction (I omit the proof).

Proposition 17. *Let* $K_1, K_2 \in \mathbb{K}$ *and define* $\dot{-}_{C_{K_2}}$ *by setting* $K_1 \dot{-}_{C_{K_2}} \alpha = K_1 \cap K_2$ *for all* $\alpha \in \mathcal{L}$. $\dot{-}_{C_{K_2}}$ *satisfies* $(\dot{-}1)$, $(\dot{-}2)$, $(\dot{-}6)$, $(\dot{-}7)$, *and* $(\dot{-}8)$*; however,* $(\dot{-}3), (\dot{-}4)$, *and* $(\dot{-}5)$ *are not satisfied.*

4. Discussion

If rational contraction were restricted to AGM partial meet contraction, unanimity or consensus as shared agreement would *not* always be available to serve as a non-question-begging position at the outset of inquiry. An agent may have no rational recourse but to beg questions against certain other parties or points of view. Whether this is (further) evidence against AGM or against the importance of unanimous consensus requires further argumentation. But there are a few concerns that should still be addressed, even if briefly. These concerns represent ways to minimize the interest of the above results.

First, one might complain about the overly conservative nature of unanimity. For approaches that attempt to assimilate reaching an epistemic consensus to voting, for instance, the presence of a single dissenter would seem insufficient grounds for excluding an otherwise unanimously accepted proposition from the consensus corpus. If complete unanimity is required for consensus, consensus is exceedingly rare, one might object. At any rate, many social and political decisions get made without it. Though I have not undertaken a defense of unanimity here, in order for the central tension of the present note to be of real interest, the unanimity conception of consensus should not be a non-starter. Whether unanimity is a nonstarter depends on the relevant *function* consensus is supposed to play. If that function is to avoid begging questions or to serve as a non-controversial basis for joint inquiry and deliberation—as at least it sometimes is—it is voting accounts, not unanimity, that may fail to get off the line.

Second, we might explore a more general notion of "accessibility" within the AGM paradigm. For example, we could ask about which belief states are accessible under some finite sequence of contractions instead of by just a single contraction. We cannot gloss over the notorious problem of iterated revision here (e.g., Spohn, 1988; Boutilier, 1993; Gärdenfors and Rott, 1995; Darwiche and Pearl, 1997; Hansson, 2003; Nayak et al., 2003). Briefly put, the problem is that AGM belief revision theory and a number of variant belief revision theories constrain only a *single* stage of belief revision. There is no account of iterated belief change in the classic AGM framework or in many of its relatives. And while various attempts to remedy this have been made (e.g., Darwiche and Pearl, 1997; Spohn, 2012), these attempts are not without substantial controversy. A compelling solution to the problem, Hansson observes, "has turned out to be very difficult to achieve" (2003, p. 42). For a recent overview and critical assessment of attempts to solve the problem, see (Booth and Chandler, 2016). And then there are others, like Levi, who deny that there is a serious problem of iterated belief change because they deny that there are diachronic norms of rationality (e.g., Levi, 1980, pp. 9 - 13). On such views, we cannot make general claims about rational, iterated belief change. So, an account of reaching consensus *via* iterated contraction awaits a compelling account of iterated contraction which so far has proved difficult to articulate and which some think is not in the offing.

More optimistically, we might think of single-shot accessibility as providing a criterion that determines with which belief states it is rational for an agent to seek consensus. While contracting to the unanimous consensus with another belief state is not possible in general, with respect to some belief states it is. In other ways, this is a pessimistic view, as it surrenders the general possibility of non-question-begging joint inquiry and deliberation.

Third, we might distinguish standards for genuinely contracting beliefs from those for hypothetically doing so (e.g., Fuhrmann and Hansson, 1994; Levi, 1996). In supposing α for the sake of the argument, an agent need not actually come to believe α. Similarly, in hypothetically contracting α, an agent need not genuinely surrender the belief. Keeping a firm grasp on the distinction between actual beliefs and hypothetical suppositions or contractions is as important in decision making as it is in epistemology. In decision making, hypothetical reasoning is used to plan for contingencies or consider the effects of potential choices. If standards of rational contraction are uniform across actual and hypothetical revision, the observations above point to a tension for AGM and the general possibility of shifting to consensus, whether we understand seeking consensus as a matter of genuinely or hypothetically revising beliefs. If different standards of contraction are appropriate depending on whether the belief change is genuine or hypothetical, AGM and universal availability of unanimity can't both be constraints on a given domain. Given that AGM provides no account of how the input sentence came to be the input sentence (that is, no full account of justified expansion or contraction), we might be tempted to think of AGM as a candidate account of hypothetical belief revision. In hypothetical belief revision, it may be reasonably maintained that there is no need to justify the input (Levi, 1996, pp. 6-7). Similarly, we might not be prepared to endorse a genuine shift to shared agreements. But hypothetical shifts to unanimous agreement are more compelling.[3] However, both the AGM axioms and universal availability of unanimity cannot be constraints on hypothetical revision.

Fourth, while AGM is indeed a very prominent account of belief revision, there are well-explored alternatives. After all, Levi, who advocates consensus as shared agreement, rejects AGM partial meet contraction in favor of a more general account that he calls *mild contraction* (Levi, 2004) *alias severe withdrawal* (Rott and Pagnucco, 1999). As we will see in the next chapter, an analogous result holds for mild contraction. Unanimous consensus is *not* always available even with mild contraction. My concern here, though, is with the tension between the AGM account of contraction and universal availability of consensus, not the details of Levi's views. However, there *are* accommodating generalizations of the AGM framework. It is shown in the appendix that package

[3]We need make no claim about whether K_1 represents the totality of beliefs or just those beliefs relevant to addressing a particular question. For example, K_1 might be the deductive closure of a particular collection of beliefs regarding some physical domain—a *theory*—aspects of which are under dispute between K_1 and K_2. In that case, urging consensus as shared agreement *on the relevant issues* would appear to be less immediately objectionable than urging such a conservative consensus position on the totality of beliefs. Agent 1 may only be concerned about consensus on the physical theory because agent 2 is an "epistemic peer" in that domain, for exmaple, and rest content with lingering disagreement on other topics. If an agent's whole view about the world is the relevant issue, the account of consensus is general and covers that case, too.

contraction, for example, allows us to state a straightforward and universally sufficient condition for shifting to shared agreements (Proposition 28). Fuhrmann and Hansson ask, "can it ever be rational to engage in [package] contraction?" Seeking consensus may provide a compelling rationale.

CHAPTER 6

Consensus Does Not Justify Contraction

1. Justifying Changes in Points of View

A model of intelligently conducted inquiry requires an account of rational changes in points of view, of replacing doubt with belief and of coming to doubt what once was believed. According to the belief-doubt model of inquiry, it is just such *changes* in beliefs—and not beliefs themselves—that stand in need of justification. Despite claims to the contrary, the difference between justifying beliefs and justifying changes in belief is substantial. Wolfgang Spohn, for instance, claims that there is no real difference since we could say that current beliefs are justified if they were acquired by some justified change, and changes are justified if the acquired beliefs are (Spohn, 2012, p. 118). Spohn's suggestion obscures the fundamental pragmatist point, reverting focus to the *pedigree* of beliefs. If changes inherit justification from acquired beliefs, the backward-facing task of justifying beliefs returns. Advocates of the belief-doubt model deny that current beliefs require justification. Coherence accounts of justification are untenable, and "opposition to foundationalism ought to be the philosophical equivalent of resistance to sin" (Levi, 2006, p. 136).

Instead, the focus in epistemology should be forward-looking, attending to questions and problems as they emerge. An agent's set of beliefs is an important resource in these efforts. Judgments of probability, truth, value, and the like are made from a certain point of view, against a background of information and commitments taken to be settled. An agent's set of *full* beliefs serves as her *standard of serious possibility.* Such a standard is crucial in inquiry and deliberation. Let $\mathcal{L}$ be a propositional language closed under truth functional operations. Belief sets are represented

We would like to thank Robby Finley, Arthur Heller, Michael Nielsen, and Ignacio Ojea Quintana for comments on earlier drafts of this paper.

by deductively closed sets of sentences of $\mathcal{L}$ (we denote a belief set K and the set of all belief sets of $\mathcal{L}$ as $\mathbb{K}$). If $\alpha \in \mathcal{L}$ is logically inconsistent with the agent's full beliefs, then α is not a serious possibility. Otherwise, α is a serious possibility. Judgments of epistemic possibility serve several important functions. The space of serious possibilities provide a measurable space over which subjective probabilities can be defined (e.g., de Finetti, 1990, p. 25). Relatedly, the state and option spaces in decision problems are determined by an agent's standard for serious possibility. Judgments of serious possibility also be taken to form the basis of a unified epistemic semantics for conditionals. On such an account, the acceptability of conditionals depends on judgments of serious possibility relative to *hypothetically* transformed states of full belief.[1]

Since *Gambling with Truth*, Levi has urged an assimilation of theoretical inquiry to practical decision making (1967a). Just as for other forms of decision making, there are *goals* relevant to changes in belief. Saying as much is not to claim that there are no differences between cognitive decision making and other sorts of decision making. The goals typical of cognitive decision making differ from those typical of, e.g., aesthetic, moral, or political choices. While the goal most relevant to a firm deliberating about a problem in production might be maximization of (expected) profit, the objective in inquiry is to maximize (expected) *epistemic* utility. Refining James' view, Levi contends that seeking valuable information and avoiding error are the aims shared by all genuine inquiries.

In the present essay, my concern is with *contraction*, the operation of giving up beliefs. There is the question of *whether* to contract, and the question of *how* to contract. At first blush, there might seem to be something of a mystery about the motivations for contraction. Since full beliefs are taken to be settled, there is no risk of importing error in contraction because nothing is imported. So the goal of avoiding error is trivially satisfied. But in contraction, information is surrendered. Why,

[1]For example, a counterfactual conditional of the form "if α were the case, then β would be the case" is evaluated by transforming (contracting) K so that α is consistent with K, then checking to see if β is in the belief set resulting from adding α to the transformed belief state. This approach to conditionals (Levi, 1996) attempts to extend Ramsey's famous acceptability condition (Ramsey, 1990a), providing epistemic semantics for a larger class of conditionals.

then, should one contract if the goals motivating belief change are avoiding error and acquiring valuable information? Levi provides two reasons. First, inadvertent expansion into inconsistency provides adequate motivation for contraction because the inconsistent belief state fails to usefully function as a standard of serious possibility. Second, whereas the concern to avoid error is "myopic" in the sense that an agent should be concerned to avoid error just at the next step of revision (and not at subsequent ones), the concern for informational value enjoys more foresight. The concern to avoid error is myopic because truth is judged, as Quine puts it, relative to our evolving doctrine "as earnestly and absolutely as can be: subject to correction, but that goes without saying" (2013). The standard for distinguishing true from false evolves with the doctrine. But as long as an inquirer is pursuing an answer to a given question, judgments of informational value will be more stable. So another motivation for contraction is to give a fair hearing to hypotheses and theories of considerable informational value. That is, the goal of obtaining valuable information at some further stage of revision can motivate contraction. Once the decision to contract has been made, the objective is to minimize the loss of informational value. We expound this point below (particularly in Section 4).

Not all changes in states of belief are deliberate. Inquirers may adopt certain *routines* for belief change. For instance, an inquirer may be (pre)committed to incorporating perceptual input, or testimony from sources she regards as reliable, etc. (A precommitment to a contraction strategy is Levi's solution for contracting from "epistemic hell," the inconsistent belief state. While deliberate expansion into inconsistency is never rational, admitting routine expansion into one's stock of belief change strategies introduces the possibility of expansion into inconsistency.) An important point here is that the inquirer is committed to incorporating the input of some testimony *before* discovering the particular content of the testimony. Routines for expansion should not be thought of as devices of marginal interest or application.

The distinction between genuine and hypothetical (or *for the sake of the argument*) belief revision is important for the argument in this essay. According to the belief-doubt model, if an agent revises to consensus between herself and another agent, such a shift should be rationalized as maximal or admissible relative to the agent's goals. Levi, however, denies that the goals characteristic of inquiry, namely, avoiding error and obtaining valuable information, are operative in seeking consensus. In Section 2, I present a challenge to justify consensus-based belief change. The distinction between hypothetical and genuine belief change is the key to responding to this challenge, as we explain in Section 3 and elaborate in Section 4.

2. Justifying a Shift to Consensus?

Much inquiry and deliberation is collective. Substantial scholarly effort has been put into studying rational collective attitudes in social choice and preference (e.g., Arrow, 1951), in probability (e.g., Genest and Zidek, 1986), and in judgment (e.g., List and Pettit, 2011). Levi writes, "conceptualizing consensus has to be a centerpiece of *any* articulate account of the intelligent conduct of problem-solving inquiry" (Levi, 2008, p. 211, emphasis ours). But it is important to point out that consensus is not just a concern in *social* epistemology and decision making. A single agent may have multiple goals or values, may entertain various rival theories, beliefs, or ways of evaluating events with respect to subjective probability. In each case, an agent may seek a neutral point of view—reasonably thought of as a consensus—from which to consider the various candidates without begging questions.

For full belief, Levi proposes the following conceptualization of consensus. Let K_1 and K_2 be two deductively closed sets of sentences. The *consensus as shared agreeement* is given by $K_1 \cap K_2$, or, for any finite set $N = \{1, 2, ..., n\}$, $K_C = \bigcap_{i \in N} K_i$. K_C is a contraction of each belief state in N. That is, $K_C \subseteq K_i$ for any $i \in N$. In shifting to K_C, no corpus runs the risk of importing error since nothing is imported into any of the corpora. But Levi also claims that informational value is not a concern. He writes, "Sometimes the aim in contraction might be not to minimize loss of informational value but to identify the shared agreements between the agent and some other agent or group of agents [...] Consensus contraction is not based on assessments of informational value but on the contents of a reference corpus" (1996, p. 27). Seeking informational value can license contraction in cases in which the agent wishes to give a hearing to an informationally valuable theory or hypothesis. This, Levi claims here, is not what drives consensus-based revision. Indeed, there may be no new, informationally valuable theory at hand for the agent to consider when contemplating the consensus position. Furthermore, in contracting, shifting to consensus does not in general minimize the loss of informational value. This point is made more precise in Section 4.

The shared agreements between some points of view constitute a noncontroversial basis for subsequent deliberation and inquiry. By restricting the assumed background to shared agreements, parties to the consensus avoid begging relevant questions against each other. But other belief states avoid begging questions, too. Any subset of the "shared agreements" belief set begs no questions. And since informational value plays no role in consensus contraction according to Levi, we might wonder what rules out strict subsets of $K_1 \cap K_2$ as legitimate consensus positions for the two belief states. Similarly, if an agent is motivated entirely by agreement with K_2 in seeking consensus, we might wonder why the agent should not simply shift to K_2. Or if K_1 is opinionated about α while K_2 is not, one might complain about calling α a *disagreement* for K_1 and K_2.

Briefly, then, the primary challenge is this. Levi makes justifying changes in points of view the focus of his epistemological outlook. The justification is decision-theoretic: changes must be optimal or admissible relative to the goals of the agent changing her mind. In changing points of view, those goals are epistemic, namely, acquiring valuable information and avoiding error. These goals, according to Levi, are not the ones guiding consensus-seeking. So what justifies shifting to a consensus in general or consensus as shared agreement in particular? Levi has not argued that consensus as shared agreement is optimal or admissible relative to certain goals.

3. Response to the Challenge

The most promising response to the challenge is simply that consensus-based belief change, when rational, is generally pretend, not genuine, belief change. There are at least three things to say about why this should be so. First, if an agent's set of full beliefs determines her standard for serious possibility, it also determines her *deliberate* response to disagreement. Suppose $\alpha \in K_1$ but $\neg\alpha \in K_2$. If K_1 is taken as the standard of serious possibility in the sense that what is inconsistent with K_1 is not seriously possible, then $\neg\alpha$ is not a serious possibility and K_2 is judged to be in error from the perspective of K_1.[2] The agent may have reasons to open up her mind concerning α, that is, to contract K_1, but those are the reasons for contraction in any case: to give a fair hearing to a hypothesis of considerable informational value or to retreat from epistemic hell. Learning that $\neg\alpha \in K_2$ may result in *routinely* importing $\neg\alpha$ into K_1, provided the agent employs some *routine* for adopting the deliverances of K_2 (for example, if K_2 is regarded as expert testimony on a relevant topic or a reliable source of information more generally), thereby transforming K_1 into the inconsistent belief state requiring a subsequent routine contraction. But the question of when to open up one's mind is not so trivial a matter as to be answered, "Automatically." Coming to doubt, too, demands justification.

Second, if it is right that neither concern to avoid error nor concern for informational value plays a role in seeking consensus, then, given the thesis that those two aims govern rational inquiry, seeking consensus is not rational inquiry. But rational inquiry is the intelligent change in point of view. So either consensus-based belief change is not rational, or it is not a matter of genuine change in point of view. The response to the constructive challenge, then, is not to articulate a justification for shifting to $K_1 \cap K_2$, but to deny that such a shift is genuine and in need of justification. From the fact that the twin aims of inquiry are not operative in seeking consensus, we infer that seeking

[2] This view may be at odds with certain "conciliatory" positions on the topic of peer disagreement, at least when sharing the same evidence is not thought of as $K_1 = K_2$, in which case, trivially, there is no disagreement.

consensus is not inquiry, not about actually or genuinely changing one's mind, even if consensus is relevant to inquiry in certain other ways.

Third, and relatedly, revising to consensus as shared agreement violates the rationality postulates that Levi endorses for belief change. Besides insufficiently motivating an affirmative answer to the question of *whether* to contract, seeking consensus can run afoul of the norms governing *how* to contract. Apparently, hypothetical belief revision need not abide by the same canons of rationality as genuine revision. The formal criteria for contraction are motivated in terms of philosophical considerations of informational value, so the point here is not entirely distinct from the point in the paragraph above (though the axioms admit of alternative motivations as Rott and Pagnucco show (1999)). We confine elaboration and demonstration of this point to the following section because the details threaten to overwhelm the idea.

Unlike genuine belief change, hypothetical revision does not require justification (Levi, 1996, p. 5). After all, the agent is not *actually* changing belief states, and, as Levi often repeats, changes are what require justifications on the belief-doubt model. Here, Levi is elaborating a pragmatist theme initiated by Peirce. Peirce claims that "the mere putting of a proposition into the interrogative form does not stimulate the mind to any struggle after belief" (Peirce, 1992a). Just as a feigned belief does not demand justification ("Suppose that α..."), "fictitious" doubt does not demand justification either.[3]

4. Levi on Contraction

In this section, I substantiate my claim above that shifting to consensus as shared agreement in general violates constraints that Levi endorses for genuine contraction.

[3]Even if consensus as shared agreement does not demand justification, there is the outstanding issue of why consensus should be as Levi identifies it. Why not, for example, take a subset of the intersection? And what properties ought the consensus position have? If consensus-based belief change is not genuine, though, the issue may be less pressing.

4.1. Preliminaries. As in the previous chapter, $\mathcal{L}$ is a language closed under truth-functional operations. Lower case Greek letters, $\alpha, \beta, ...$ (except γ, which is reserved for another use), refer to sentences of $\mathcal{L}$, and capital Roman letters, $A, B, ...$, refer to sets of sentences. Atomic formulae are denoted by lower case Roman letters, $p, q, ...$ $Cn : \mathscr{P}(\mathcal{L}) \to \mathscr{P}(\mathcal{L})$ is a Tarskian consequence operator. As is standard, we suppose that Cn satisfies the properties listed in the previous chapter. A set of sentences K is called a *theory* if $Cn(K) = K$, i.e., K is deductively closed. Let $\mathbb{K}$ denote the set of all deductively closed subsets of $\mathcal{L}$. Elements of $\mathbb{K}$ are *belief sets*, with sentences representing beliefs. We call K a *maximal consistent* set if for any $\alpha \in \mathcal{L}, \alpha \in K$ or $\neg\alpha \in K$, but not both. Let $\mathcal{L}\perp^{\perp}$ denote the set of maximal consistent sets of $\mathcal{L}$.

In the AGM theory of belief revision, there are three standard operations of belief change (e.g., Alchourrón et al., 1985; Gärdenfors, 1988). *Expansion* adds beliefs to K. *Revision* adds beliefs in such a way that the result is a consistent theory (so long as the input sentence is not itself a contradiction). *Contraction* removes beliefs from K. Since we are dealing with theories, how to remove some sentence from K is not uniquely determined. In general, contraction requires a choice among subsets of K that fail to imply the sentence to be removed. The AGM contraction postulates are as follows.

($\dot{-}$1) $K \dot{-} \alpha = Cn(K \dot{-} \alpha)$ whenever $K = Cn(K)$

($\dot{-}$2) $K \dot{-} \alpha \subseteq K$

($\dot{-}$3) If $\alpha \notin K$, then $K \subseteq K \dot{-} \alpha$

($\dot{-}$4) If $\alpha \in K \dot{-} \alpha$, then $\alpha \in Cn(\varnothing)$

($\dot{-}$5) $K \subseteq Cn((K \dot{-} \alpha) \cup \{\alpha\})$

($\dot{-}$6) If $Cn(\alpha) = Cn(\beta)$, then $K \dot{-} \alpha = K \dot{-} \beta$

($\dot{-}$7) $K \dot{-} \alpha \cap K \dot{-} \beta \subseteq K \dot{-} (\alpha \wedge \beta)$

($\dot{-}$8) If $\alpha \notin K \dot{-} (\alpha \wedge \beta)$, then $K \dot{-} (\alpha \wedge \beta) \subseteq K \dot{-} \alpha$

The first six are called the *basic* postulates; the last two are known as the *supplementary* postulates.

Some constructions of central focus in the AGM theory of contraction are *partial meet contraction* and a construction in terms of *epistemic entrenchment orderings.* We look at a variation of partial meet contraction in characterizing Levi-contractions in the next section. Orderings of epistemic entrenchment are the key construction in the section on *severe withdrawals.*

4.2. Levi-Contraction. Levi's account of contraction has evolved over the years. Stages include undamped informational value (Levi, 1991), damped informational value (version 1) (Levi, 1991, 2004), and, finally, damped informational value (version 2) (Levi, 1996, 2004). We look first at the characterization of partial meet *saturatable* contraction due to Hansson and Olsson (1995). Other versions of Levi's account are special cases of this account (Rott and Pagnucco, 1999, pp. 28-29).[4]

Hansson and Olsson's characterization of Levi-contraction requires a new postulate.

($\dot{-}$9) If $\alpha \in Cn(\varnothing)$, then $K \subseteq K \dot{-} \alpha$

Fuhrmann and Hansson dub this postulate *failure* (1994). The stronger version of the postulate (with equality in the consequent) that is assumed in (Hansson and Olsson, 1995) follows from ($\dot{-}$2) and ($\dot{-}$9). ($\dot{-}$9) follows from ($\dot{-}$1) and ($\dot{-}$5), but Levi's account of contraction does not generally satisfy ($\dot{-}$5), the controversial postulate of *recovery* (Hansson, 1999, pp. 86-87).

A set of sentences $K \subseteq \mathcal{L}$ is *α-saturatable* if and only if $K = Cn(K)$ and $Cn(K \cup \{\neg\alpha\})$ is a maximal consistent set. The set of all α-saturatable subsets of a theory plays an essential role in Levi-contraction.

[4]We note, however, that Hansson and Olsson's treatment of Levi-contraction may fail to capture all of the philosophical subtitles of Levi's account. In particular, Levi appeals to a fairly intricate semantics for contraction in terms of partitions.

DEFINITION 2. Let K be a theory in $\mathcal{L}$ and let $\alpha \in \mathcal{L}$. The set of α-saturatable subsets of K is given by

$$S(K, \alpha) = \{K' \subseteq K : (i)\ K' = Cn(K') \text{ and } (ii)\ Cn(K' \cup \{\neg\alpha\}) \in \mathcal{L}\bot^{\bot}\}.$$

Levi-contractions can be defined in terms of selection functions on sets of saturatable subsets of K. A correspondence $\gamma : \mathscr{P}(\mathbb{K}) \rightarrow \mathscr{P}(\mathbb{K})$ is a selection function on $S(K, \alpha)$ such that, for all α, if $S(K, \alpha) \neq \varnothing$, then: (i) $\gamma\big(S(K, \alpha)\big) \subseteq S(K, \alpha)$ and (ii) $\gamma(S(K, \alpha)) \neq \varnothing$. When $S(K, \alpha) = \varnothing, \gamma(S(K, \alpha)) = \{K\}$.

DEFINITION 3. A function $\dot{-}$ is a Levi-contraction for K just in case there exists a selection function γ such that for all $\alpha \in \mathcal{L}$:

$$K \dot{-} \alpha = \begin{cases} \bigcap \gamma\big(S(K, \alpha)\big), \text{ if } \alpha \in K \\ K, \text{ otherwise} \end{cases}$$

THEOREM 15. *(Hansson and Olsson, 1995, Thm. 4.1)* $\dot{-}$ *is a Levi-contraction for* K *if and only if* $\dot{-}$ *satisfies* $(\dot{-}1), (\dot{-}2), (\dot{-}3), (\dot{-}4), (\dot{-}6)$, *and* $(\dot{-}9)$.

In the previous chapter, it was shown that consensus as shared agreement is not reachable for all AGM partial meet contraction operators on a belief set, and that, for certain belief sets, consensus is not available for *any* partial meet contraction operator on those sets (Propositions 15 and 16). Here, we prove propositions in that same vein, showing that there are consensus contractions of K_1 that, *for a fixed contraction operator*, are *not* Levi-contractions or mild contractions. If Levi-contraction or mild contraction is supposed to account for all rational contractions, the upshot of the propositions to come is that genuine consensus contraction is not generally rational.

PROPOSITION 18. *It is* ***not*** *the case that, for any belief sets* $K_1, K_2 \in \mathbb{K}$ *and any Levi-contraction,* $\dot{-}$*, there is a sentence* α *such that* $K_1 \dot{-} \alpha = K_1 \cap K_2$.

4.3. Mild Contraction *alias* Severe Withdrawal. Damped informational value (version 2) contraction *alias* mild contraction *alias* severe withdrawal has been axiomatized by Rott and Pagnucco (1999) and Arló Costa and Levi (2006). The axioms include those for Levi-contractions and two additional axioms: ($\dot{-}8$) above and a strengthening of ($\dot{-}7$).

($\dot{-}7a$) If $\alpha \notin Cn(\varnothing)$, then $K \dot{-} \alpha \subseteq K \dot{-} (\alpha \wedge \beta)$

Postulate ($\dot{-}7a$) is known as *antitony* (short for "anti-monotony") in the literature, and, like recovery, finds ready opposition. For example, Hansson claims that antitony without the hedge in the antecedent "does not hold for any sensible operator of contraction" (1999, p. 117).

Rott follows Makinson in distinguishing between contraction and withdrawal functions. Withdrawal functions are those operations of belief removal satisfying ($\dot{-}1$), ($\dot{-}2$), ($\dot{-}3$), ($\dot{-}4$), and ($\dot{-}6$). Hansson and Olsson's Levi-contractions, then, are those withdrawal functions satisfying failure. AGM contraction operators are those withdrawal functions that satisfy recovery. I, however, will follow the convention of calling any operation removing beliefs a *contraction*. On a decision-theoretic rendering of the problem of contraction, when the concern is to remove α, $K \dot{-} \alpha$ should be optimal or admissible relative to the set of all available options. Artificially restricting the option set to what Makinson calls contractions misconstrues the nature of the decision problem (2004). In light of Theorem 15, inspecting the Levi-contraction postulates reveals that, unlike the case with AGM partial meet contraction, all the possible ways to "withdraw" α from K are representable by meets of α-saturatable subsets of K (Rott and Pagnucco, 1999, p. 531).

Rott and Pagnucco's characterization of mild contraction/severe withdrawal is in terms of *epistemic entrenchment orderings*. Since entrenchment orderings are familiar constructions from the literature on belief revision theory, we deal with them first. Levi's semantics for mild contraction will be discussed afterwards. An entrenchment ordering of sentences of $\mathcal{L}$ is intended to represent epistemic value, or usefulness in inquiry and deliberation, or vulnerability to be given up in contraction and revision.

DEFINITION 4. Let $\leqslant$ be a binary relation on $\mathcal{L}$. We call $\leqslant$ an *epistemic entrenchment ordering* with respect to some $K \in \mathbb{K}$ if the following conditions are satisfied:

$(E1)$ If $\alpha \leqslant \beta$ and $\beta \leqslant \delta$, then $\alpha \leqslant \delta$

$(E2)$ If $\beta \in Cn(\{\alpha\})$, then $\alpha \leqslant \beta$

$(E3)$ $\alpha \leqslant \alpha \wedge \beta$ or $\beta \leqslant \alpha \wedge \beta$

$(E4)$ If $K \neq \mathcal{L}$, then: $\alpha \leqslant \beta$ for every $\beta \in \mathcal{L}$ iff $\alpha \notin K$

$(E5)$ If $\beta \leqslant \alpha$ for every $\beta \in \mathcal{L}$, then $\alpha \in Cn(\varnothing)$.

It follows, as Rott and Pagnucco remark, that $\leqslant$ is a weak ordering of (a complete and transitive binary relation on) $\mathcal{L}$ with tautologies as maximal elements and non-beliefs as minimal elements (1999, p. 526). The claims about the minimal and maximal elements of $\leqslant$ follow immediately from $(E4)$ and $(E5)$, respectively. $(E1)$ is itself a statement of the transitivity of $\leqslant$. Only the completeness of $\leqslant$ remains. Let $\alpha, \beta \in \mathcal{L}$. By $(E3)$, either $\alpha \leqslant \alpha \wedge \beta$ or $\beta \leqslant \alpha \wedge \beta$. Suppose, without loss of generality, the former. By $(E2)$, we also have $\alpha \wedge \beta \leqslant \beta$ since $\beta \in Cn(\{\alpha \wedge \beta\})$. An application of $(E1)$ yields $\alpha \leqslant \beta$. So, $\leqslant$ is complete.

A contraction operator can be recovered from a given entrenchment ordering. Rott's recipe is the following.

DEFINITION 5. $K \dot{-} \alpha = \begin{cases} K \cap \{\beta : \alpha < \beta\}, & \text{if } \alpha \in K \text{ and } \alpha \notin Cn(\varnothing) \\ K, & \text{otherwise} \end{cases}$

Similarly, an entrenchment ordering can be constructed from a given contraction function.

DEFINITION 6. $\alpha \leqslant \beta$ iff $\alpha \notin K \dot{-} \beta$ or $\beta \in Cn(\varnothing)$

Entrenchment-based contraction operators (in the sense of Definitions 4 and 5) are characterized by the severe withdrawal/mild contraction postulates.

Theorem 16. *(Rott and Pagnucco, 1999, Observation 19)* $\dot{-}$ *is an entrenchment-based severe withdrawal (mild contraction) function for* K *if and only if* $\dot{-}$ *satisfies* ($\dot{-}$1), ($\dot{-}$2), ($\dot{-}$3), ($\dot{-}$4), ($\dot{-}$6), ($\dot{-}$7*a*),($\dot{-}$8), *and* ($\dot{-}$9).

It is important to note that what Levi refers to as *mild* and Rott and Pagnucco refer to as *severe* are two different things. In the latter case, it is the loss of *information* incurred by the operation of mild contraction/severe withdrawal (measured by subset inclusion) that is severe. Levi, however, has long argued that the focus in contraction should be on minimizing the loss of *informational value* and not information *tout court*. To this end, Levi introduces a *value function*, $V : \mathbb{K} \to \mathbb{R}$, and argues that, minimally, V should satisfy the following principle.

Weak Monotony. For all $K, K' \in \mathbb{K}$, if $K \subseteq K'$, then $V(K) \leqslant V(K')$

A *value-based* Levi-contraction is a Levi contraction for which the selection function respects informational value in the following sense:

$$\gamma\big(S(K,\alpha)\big) = \{K' \in S(K,\alpha) : V(K'') \leqslant V(K') \text{ for all } K'' \in S(K,\alpha)\}$$

Value-based Levi-contractions—an intermediate between Levi-contractions and mild contractions—are characterized by Arló Costa and Liu (Arló-Costa and Liu, 2010). Levi (2004) and Arló Costa and Levi (2006) argue for the following additional principle for measures of informational value.

Weak Min. For every finite $\mathscr{K} \subseteq S(K,\alpha)$, $V(\bigcap_{K' \in \mathscr{K}} K') = min_{K' \in \mathscr{K}} V(K')$

Value-based Levi-contractions such that V satisfies the above principles are characterized by the mild contraction/severe withdrawal postulates (Arló-Costa and Levi, 2006).

However, consensus is not always one mild contraction away.

Proposition 19. *It is **not** the case that for any belief sets* $K_1, K_2 \in \mathbb{K}$ *and any mild contraction operator,* $\dot{-}$*, there is a sentence* α *such that* $K_1 \dot{-} \alpha = K_1 \cap K_2$.

4.4. Consensus-Based Contraction. Moreover, if we follow Levi's proposal to define a consensus-based contraction operator (1996, p. 27) (as considered in the previous chapter), $\dot{-}_{C_{K_2}}$ such that $K_1 \dot{-}_{C_{K_2}} \alpha = K_1 \cap K_2$ for any $\alpha \in \mathcal{L}$, it is clear that $\dot{-}_{C_{K_2}}$ is not a Levi-contraction for K_1.

PROPOSITION 20. *Let $K_1, K_2 \in \mathbb{K}$ and define $\dot{-}_{C_{K_2}}$ by setting $K_1 \dot{-}_{C_{K_2}} \alpha = K_1 \cap K_2$ for all $\alpha \in \mathcal{L}$. Then, $\dot{-}_{C_{K_2}}$ does not satisfy ($\dot{-}$3), ($\dot{-}$4), or ($\dot{-}$9).*

It follows from Proposition 20 that $\dot{-}_{C_{K_2}}$ is not a mild contraction function, a Levi-contraction, nor even a withdrawal operator in the terminology of Makinson and Rott and Pagnucco. What is most important to stress is that such a function does not satisfy the constraints imposed by minimizing the loss of informational value, and so does not abide by the principles of rational contraction.

The point I wish to make in this section is that if rational contraction is captured by (perhaps some suitably constrained version of) Levi-contraction, then genuine consensus-based belief change is not rational in general. Put another way, an agent committed to a contraction strategy represented by a certain Levi contraction operator may be unable to achieve consensus with certain other belief sets. We should understand Levi, not as advocating both an account of rational belief change and a type of belief change that is irrational on that account, but as claiming that consensus prompts belief change for the sake of the argument (e.g., Levi, 1991; Levi, 1996, p. 27), and is not a generally appropriate rationale for genuine belief change.

5. Some Objections Headed Off

The thesis of this chapter, stated a little provocatively, but nonetheless genuinely, is that *genuine* consensus-based belief change is not *rational*.

5.1. In General. One response to the case that we make in Section 4 might be that, while it may not be rational to seek consensus with just *any* other point of view, it is sometimes rational. That is, there are belief sets with respect to which consensus is available for K_1 by Levi-contraction, and, with those states, it is at least possibly rational to revise to the consensus. Levi is not concerned to deny this—in fact, he cannot—but to urge that such revision should be motivated and guided by the principles governing rational contraction in general, i.e., minimizing the loss of informational value. In such cases, the consensus position simply happens to be an admissible contraction.

5.2. Single-Sentence, One-Shot Contraction. Propositions 18 - 20 demonstrate limitations for shifting to consensus in the frameworks of Levi-contraction, mild contraction, and consensus-based contraction (as prior propositions demonstrate for AGM contraction). But these frameworks are limited to a single contraction aiming to remove a single sentence.[5] Two generalizations suggest themselves. First, we could consider frameworks that allow contracting multiple sentences (e.g., Fuhrmann and Hansson, 1994). Second, we could consider frameworks with accounts of *iterated* contractions (e.g., Darwiche and Pearl, 1997). In the setting of *package contraction*, for example, consensus as shared agreement between any K_1 and K_2 *is* always available for K_1 in a single contraction by contracting the set-theoretic difference, $K_1 \backslash K_2$ (Proposition 28). Similarly, consensus might be achieved under iterations of contraction.

The obvious first point to make is that, to the extent that the accounts of belief revision presented above are normatively compelling, alternatives are not. The arguments that support, say, mild contraction's claim to normative propriety are arguments against going about contraction another way. But I doubt that this is a fully compelling response with respect to package contraction.

Regarding iteration, it is true that AGM belief revision theory and a number of variants fail to constrain iterated belief revision. Many see this as a serious cause for concern about these theories. Again, a compelling solution to the so-called *problem of iterated revision*, Hansson observes, "has turned out to be very difficult to achieve" (2003, p. 42). Attempts to respond to the perceived difficulty are subject to substantial outstanding controversy. So, even granting the need for an

[5]Whether this is accurate of consensus-based contraction is a legitimate question since the "input" sentence is irrelevant.

account of iterated revision, the limitations indicated in the propositions in the present chapter are significant given theories of belief change on offer.

But some students of belief revision deny that there is any problem about iterated revision. In particular, Levi denies that there are any diachronic norms of rationality, any constraints for iterated revision awaiting articulation: "All conditions of rationality are equilibrium conditions. In a sense they are synchronic conditions [...] Furthermore, in stating conditions of rational equilibrium, no prescription is made regarding the psychological path to be taken in moving from disequilibrium or from one equilibrium position to another. In other words, there are no norms prescribing rational learning processes" (Levi, 1970, p. 137). Despite the focus on justifying *changes* in points of view, rationality constraints are not themselves diachronic on the belief-doubt model. The *comparative statics* view expressed in the quotation above sees sets of commitments—epistemic and valuational—as equilibrium states determining rational courses of action in inquiry and deliberation (see also Levi, 1980, pp. 9-13).

5.3. Rational Revision Outside of Inquiry. There is a thin sense of rationality according to which consensus-based belief revision could be rational. An action—including a revision of belief—is rational if it is optimal or admissible relative to the agent's goals. In inquiry, those goals are avoiding error and acquiring valuable information. But shifting to the consensus position could be instrumentally rational relative to *some* end or set of ends. Various kinds of questions and problems arise occasioning various forms of inquiry to address them. Such inquiries have potentially very diverse goals. The plausibility of the contention that consensus-based revision is not rational depends, in part, on the thesis that seeking valuable information and avoiding error are the (proper) goals common to all inquiries. That is a substantive thesis outstripping the minimal constraints articulated in theories of rational choice.

It is also a thesis not without detractors. Even within the pragmatist tradition, Dewey, for example, seems at points *not* to have taken truth to be a goal in inquiry (Dewey, 1941), and Rorty claims to be following Dewey in such disregard in his interpretation of warranted assertability (Rorty, 1980; Levi, 1998). No proof, it must be admitted, can be adduced for such a thesis about the legitimate goals common to all inquiries. As is often the case in philosophy, the best that can be offered is a comparison of the fruits of this view with those of other views. However, I will not undertake such a comparison at this point.

5.4. Group Agents. In the social choice, probability pooling, and judgment aggregation literatures, the aggregate opinion has been interpreted in various ways. The aggregate view might be thought of as a compromise adopted for the purposes of performing some exercise in group inquiry or deliberation (though no group member may genuinely change her view), or as the "rational" opinion to adopt on being informed of the opinions of the (other) group members. A rather distinct interpretation has it that the group can constitute an agent, subject to the same or similar rationality constraints as other sorts of agents (Rovane, 1997; List and Pettit, 2011). On such views, aggregation may yield the opinion of the group agent. Doesn't the group opinion need to be justified?

Suppose that the belief set of the group agent, G, is given by the meet of the belief sets of the group members: $K_G = \bigcap_{i \in N} K_i$. What belief *change* has G undergone? G is not shifting from K_i (for some $i \in N$) to K_G. Rather, K_G is constituted by the meet of the K_i; G is not changing points of view. Since only *changes* in belief require justification, the question of justification does not yet arise for G. When it does arise, it should be handled analogously to the issue of justification for any other agent changing points of view.

6. Conclusion

Consensus is important, but it is not the impetus for genuine, rational belief change.

CHAPTER 7

Social Choice Theory for Deliberative Democrats

1. Introduction

Social choice theory and deliberative democracy are the two most prominent traditions in contemporary democratic theory (Dryzek and List, 2003; Estlund, 2005). Social choice theory, as studied in economics (Sen, 1997), political science (Elster, 2007), and other social sciences, concerns relationships among the preferences of *individuals* and *collective*, *social*, or *group* preferences and choices. In its modern form, social choice theory was initiated by Kenneth Arrow's Impossibility Theorem, which shows that a small number of minimal and compelling constraints cannot be jointly satisfied by quite a broad class of methods for aggregating individuals' preferences into a social preference ordering of the options (1951). Given that *voting* is among such methods, many have found this result disturbing. Amartya Sen observes that impossibility results in social choice theory "have often been interpreted as being thoroughly destructive of the possibility of reasoned and democratic social choice" (2004a). For instance, Russell Hardin claims that, from social choice theory, we have learned of "flaws—grievous foundational flaws—in democratic thought and practice" (1993).

According to deliberative democrats, the democratically legitimate way to form social preferences or make collective decisions is by public deliberation and reasoning among the members of the social group (Cohen, 1997). In deliberation, members discuss the relevant issues and are open to reason-based (not coercion-based) revision of their preferences. The aims and foundations of deliberative democracy, claims William Riker, face very severe challenges from social choice theory

(1982). Social choice theory, on this view, "demonstrates the impossibility, instability or meaninglessness of the rational collective outcomes" pursued in deliberative democracy (Dryzek and List, 2003, p. 2).

This assessment of the state of affairs is hasty. In what follows, I offer a reconciliation. A generalization of the standard notion of preference aggregation is presented. The mathematical framework of preference aggregation is generalized by employing sets of preference orderings to represent social preference (that is, I propose using *set-valued* aggregation functions). I study some particular models of preference aggregation that allow for *indeterminacy* in social preferences as a proof of concept. I argue that these models have stronger *normative* motivations than standard preference aggregation models and deliberative democrats have reasons to endorse them. In rough outline, aggregation produces a *set* of socially permissible orderings; deliberation, not aggregation, may then be invoked as a means of strengthening consensus or decreasing indeterminacy in social preference. For these models, possibility results can be proved, showing that certain sets of axioms of central concern in social choice theory—namely, natural extensions of those of relevance to Arrow's result and those involved in Sen's Impossibility of the Paretian Liberal—are jointly satisfiable (Section 6). As is typical in social choice theory, however, other formal and conceptual limitations emerge (Section 9.1). I argue that the importance of the latter sort of results has been overemphasized, especially from the perspective of deliberative democracy and, more generally, conceptions of consensus akin to the one presented below.

The essay is structured as follows. After presenting the basic framework of social choice theory and some of the central limitative results (Section 2), I very briefly canvass some attempts to get around such negative results in order to situate the project pursued here (Section 3). Section 4 introduces and offers philosophical motivation for the account of preference aggregation studied in this chapter. Over the course of the next two sections, I develop a generalization of the canonical social choice framework, framing aggregation with indeterminate preferences in the mathematical language common in work on social choice. I employ *set-valued* aggregation functions and extend popular aggregation axioms accordingly. Two concrete examples of the format are considered, and *possibility* results are established: the constraints relevant for Arrow's and Sen's results can be simultaneously satisfied. While some of the technical details of the following sections are tedious, they are in service of some important conceptual points. Some of the model's key payoffs for social and political theory are discussed in Sections 7 and 8. In particular, I offer resolutions of some important conceptual tensions in democratic theory and welfare economics. Section 9 confronts some challenges, including, importantly, some of the limitative results that are not generalized away (9.1), and addresses the crucial issue of *using* a set of preference orderings to make choices (9.2).

2. Social Choice

Let X be a universal set of alternatives or options. The options could be business plans, political candidates, plans for dinner, public health policies, research agendas, whatever. Let N be a set of individuals, labeled $i = 1, ..., n$. $R \subseteq X \times X$ is a *preference relation* if R is transitive and complete, i.e., a weak order.[1] xRy is read "x is (weakly) preferred to y." I will use the infix and graph notations interchangeably. The set of all weak orderings of X is denoted $\mathfrak{R}$. To each individual, i, we associate an ordering, R_i. A profile is a tuple, $(R_1, ..., R_n) \in \mathfrak{R}^n$, of individual orderings of the alternatives. A *social welfare function* for the n individuals is a function

$$F: D \to \mathfrak{R}$$

mapping profiles of individual orderings to a social ordering. $D \subseteq \mathfrak{R}^n$ is the domain of possible profiles of individual preference orderings. R (without a subscript) will be used as shorthand for $F(R_1, ..., R_n)$, the social or group or consensus preference ordering.

[1] R is *transitive* if, for all $x, y, z \in X$, if xRy and yRz, then xRz. R is *complete* if, for all $x, y \in X$, xRy or yRx. Weak orders are consequently *reflexive*: for all $x \in X$, xRx.

In his doctoral dissertation, Arrow proved an extremely surprising result. Given the above setup, he showed that social welfare functions (SWFs) cannot satisfy a small set of very compelling conditions. One presentation of his result is as follows. The first condition says that no restrictions are placed on the structure of individual preferences. Individuals may order the options however they see fit. The domain of the aggregation function includes all possible profiles for X.

Universal Domain (U). The domain of F is the set of all possible profiles, i.e., $D = \mathfrak{R}^n$.

The next constraint has a prestigious pedigree in the history of economic thought, but we need a bit more notation for the formulation below. Any binary relation can be split into its symmetric, I, and asymmetric, P, parts. We put xPy iff xRy and $\neg yRx$ (i.e., $P = R \backslash R^{-1}$).[2] And we put xIy iff xRy and yRx. In the setting of preference relations, P can be understood as *strict* preference, and I as *indifference*. Then, weak preference, R, is composed of the union of P and I: $R = P \cup I$. The constraint requires that if every individual strictly prefers one option to another, then that agreement is reflected in the social ordering.

Pareto Unanimity (P). For all $x, y \in X$ and all profiles in D, if xP_iy for all $i \in N$, then xPy.

The following requirement demands that the social ranking of two options x and y depends only on the individual rankings of x and y and not other, "irrelevant" options.

Independence of Irrelevant Alternatives (IIA). For all $x, y \in X$ and any $(R_1, ..., R_n)$, $(R'_1, ..., R'_n) \in D$, if xR_iy iff xR'_iy for all $i \in N$, then xRy iff $xR'y$.

Finally, F should admit of no dictators of social preference. A dictator for an aggregation rule determines the social preference regardless of the profile.

Non-Dictatorship (D). There is no $i \in N$ such that, for all $x, y \in X$ and all $(R_1, ..., R_n) \in D$, if xP_iy, then xPy.

[2] $R^{-1} = \{(y, x) \in X \times X : (x, y) \in R\}$ is the *inverse* of R.

Arrow shows that, though these conditions may be quite compelling, they are not jointly satisfiable. In other words, the four axioms for F are jointly inconsistent (when the set of alternatives, X, contains more than two options).

THEOREM 17. *(Arrow, 1951) Suppose $|X| > 2$. If F satisfies **U**, **P**, and **IIA**, then there is some $i \in N$ that is a dictator for F.*

Sen similarly has shown an important limitation for aggregating preferences. He proposes a constraint intended to capture a minimal notion of social liberalism. Each person should have *some* choice over which he or she is decisive (e.g., whether to read *Lady Chatterley's Lover*, or whether to paint one's bedroom walls pink).

Liberalism (L). For every $i \in N$ there is a pair of options $x, y \in X$ such that xP_iy implies xPy, and yP_ix implies yPx.

Sen considers various ways of weakening **L** in order to prove stronger impossibility results (e.g., at least two individuals have options over which they are decisive). Such variations need not concern us here because, as will be shown, certain generalizations of this strong version of the axiom are satisfiable on the models we study below. Similarly, Sen stated his result for *social decision functions* instead of for social welfare functions. The class of social decision functions includes the class of social welfare functions, so his result applies to social welfare functions as well. We can now state Sen's Impossibility of the Paretian Liberal.

THEOREM 18. *(Sen, 1970b) No social welfare function satisfies **U**, **P**, and **L**.*

3. Ways to Cope with the Impossibility Results

A number of ways to dodge Arrow's result and related ones have been explored. Though they may be compelling initially, each constraint on F might be relaxed. Certain constraints are plainly more attractive candidates for relaxation than others. For aggregation of preference orders, ways of restricting the domain of the SWF enjoy wide appeal. That is, allowing individuals to order the options in just *any* way as mandated by **U** may be *too* general. Perhaps SWFs should only take other-regarding preferences, for example, as input. This, needless to say, is a restriction on the domain of the SWF. Still others have proposed certain structural conditions characterizing the sorts of preferences suitable for aggregation. Restriction to so-called "single-peaked" preferences finds advocates not only for the various "nice" features such preferences display (Black, 1958), but also for the potential to resolve the tension between social choice theory and deliberative democracy (Miller, 1992).

Pareto Unanimity, **P**, on the other hand, is relatively sacrosanct. Relaxing it is a much less popular route around Arrow's or Sen's results, even if not unexplored (e.g., Mongin, 1997). Perhaps even more so, Non-dictatorship (**D**) has been regarded as pretty much non-negotiable, at least for the standard applications of the social choice framework. Much more popular is relaxing Independence (**IIA**). List claims, "Almost all familiar voting methods over three or more alternatives that involve some form of preferential voting (with voters being asked to express full or partial preference orderings) violate this condition" (2013).

We can also consider changing some of the other formal features of the social choice framework. For example, the framework presented above does not allow interpersonal comparisons of utility. SWFs take profiles of only preference orderings as input. Preference orderings do not contain information about strength of preference or about how to compare different preferences with one another. According to Arrow, "interpersonal comparison of utilities has no meaning and ... there is no meaning relevant to welfare comparisons in the measurability of individual utility" (Arrow, 1951, p. 9). The "Hicksian" or "ordinalist" revolution ushered in the ban on interpersonal comparisons of utility in economic theory, but it is far from clear that it is a defensible ban and it is also far from being universally accepted today. For advocacy for "widening the informational basis" of the social choice framework beyond mere ordinal information, see for example (Harsanyi, 1953; Sen, 1977;

Levi, 1986a; Sen, 2004a). In this chapter, however, I will largely accept the restriction to ordinal information in an effort to evince the potential of indeterminacy even with a narrow informational basis. But I will indicate some important limitations that this imposes on the present study (e.g., Section 9.3).

It seems to me that most of the avenues discussed above are interesting, plausible, and worth pursuing (especially widening the informational basis). For nice overviews of these various approaches and associated results, see for example (Suzumura, 1983; List, 2013). To delimit the approach pursued here, consider a final way of altering the framework. Sen and others explore several ways of generalizing social welfare functions. For instance, we could consider what he calls *social decision functions* (SDFs) that take values in the set of relations that generate choice functions, instead of just the set of weak orders. The set of binary relations on X that generate choice functions includes the set of weak orders. I study social welfare functions generalized to *set-valued* functions. As shown below, the account does *not* reduce to binary social relations. Indeterminacy enters into the aggregation framework precisely by relaxing the requirement that the output of an aggregation function is a unique preference ordering. That is, I pursue a particular approach to relaxing the ordering assumption for social preference, the assumption that the consensus, group, or social preference ought to be characterized by a unique weak order (or even by a unique binary relation). Such a requirement is excessively restrictive for individual agents, let alone groups. In the next section, I motivate this formal maneuver.

4. Philosophical Motivation for Indeterminacy

4.1. Interpretation. It is worth pausing to reflect on some questions of interpretation. What, precisely, does the output of a social welfare function represent? $F(R_1, ..., R_n)$ might be interpreted as some sort of *rough summary* of the n preference relations. Or it could represent a *compromise* adopted by the group in order to complete a collective decision making task. In this case, perhaps we should think of each individual as facing a private decision problem, the goal of which is to ensure a social preference relation that secures the best shot at outcomes he or she privately prefers. I think that it is—or, at any rate, should be—uncontroversial that when all individual preferences are the same, $F(R_1, ..., R_n)$ admits a sensible interpretation as social *consensus*, provided that $F(R_1, ..., R_n)$ agrees with the unanimous preferences.[3] Saying more than that courts controversy.

Shortly after the publication of Arrow's work, Buchanan voiced strong interpretive criticisms of the social choice framework. The criticism takes as its jumping off point the attribution of rationality to groups. To what do standards of rationality apply? If we attempt to address that question by first considering standard examples of rationality constraints—say, consistency in beliefs or coherence in probability judgments or preferences—it appears as if canons of rationality apply to *points of view* or, by extension, things having points of view, *agents*. At any rate, such constraints are appealed to in Arrow's result and in the social choice literature more generally. The social preference relation, for instance, is required to be transitive. Indeed, these are the sorts of constraints that concern Buchanan. So, he reasons, if it makes sense to assess a collective in terms of rationality, it must make sense to think of the collective as having a point of view. In particular, it must make sense to attribute preferences to the group.

> Rationality or irrationality as an attribute of the social group implies the imputation to that group of an organic existence apart from that of its individual components. If the social group is so considered, questions may be raised relative to the wisdom or unwisdom of this organic being. But does not the very attempt to examine such rationality in terms of individual values introduce logical inconsistency at the outset? Can the rationality of the social organism be evaluated with any value ordering other than its own?
>
> The whole problem seems best considered as on of the "either-or" variety. We may adopt the philosophical bases of individualism in which the individual is the only entity possessing ends or values. In this case, no question of social or collective rationality may be raised. A social value scale as such simply does not exist. Alternatively, we may adopt some variant of the organic philosophical assumptions in which the collectivity is an independent entity possessing its own value ordering. It is legitimate to test the rationality or irrationality of this entity only against this value ordering. (Buchanan, 1954)

[3]Sometimes the distinction between the consensus and compromise interpretations is conceived in terms of the distinction between groups that form a *team*, i.e., the individuals share a common goal, and groups that do not face a collective decision problem, but n individual decision problems.

So if the group has preferences of its own, the claim is, individual preferences are irrelevant to the rationality of group decisions. If only individuals have preferences, the question of social preference and rationality does not arise. A defender of social choice theory could take either horn of Buchanan's dilemma, allowing attributions of rationality to groups or not. In the following section, I will argue that on either horn, the appropriate representation of "social" preference is generally a set of preference orderings and not a single preference ordering.

The first horn countenances group agency in the sense that groups can be reasonably assessed in terms of rationality. The practice of treating groups analogously to individuals is of ancient origin. In *The Republic*, for example, Plato suggests a systematic analogy between the organization and proper functioning of the soul and of the state. Much economic theorizing treats collectives as agents: preferences are often attributed to families and firms. Levi gives voice to a very permissive view of rational agency:

> Any system, whether it is animal, vegetable or mineral, whether it is automaton, human or a group of automata or humans, can qualify as an agent for the purpose of discussing rational choice [...] provided that choices, beliefs, preferences, values and goals are ascribable to the system and provided that is is appropriate to urge conformity to norms of rational preference, belief and choice. Saying this does not imply that all social groups act as agents or that those which do do so all the time. However, we cannot claim more for animals, automata or even human beings. (1986a, p. 153)

The crux of the matter, though, would seem to be the conditions under which preferences are "ascribable" and urging conformity to rationality norms "appropriate."

Supposing it makes sense on some occasions to assess groups in terms of rationality norms, is the group's point of view a function of individual points of view? Arrow's framework assumes so. Buchanan disputes this feature. But what is Buchanan's case? Nothing prevents one agent's preferences from being a function of another's. An agent could even assume another agent's preference ordering as her own. In the case at hand, an agent is concerned with the preferences or "welfare" of multiple agents. So, testing the "rationality or irrationality" of the group agent may require reference to the individual preference orderings after all.

One related point worth emphasizing is that aggregation has an interesting interpretation for *single agents*. A single agent may have multiple goals relevant to a decision problem, inducing conflicting orderings of the options. An *individual* may need some way of forming preferences based on more than one ordering of the options. For example, a person may value things contributing to her health like exercising and cooking nutritious meals. She may also have professional goals that require dedicating as much time as possible to her work. In such a case, assessing the rationality of the person may require reference to the various relevant preference orderings after all. Those who allow the extension of agency to groups would be quick to point out that one agent's preferences may be a function of the orderings induced by various objectives. So there would seem to be no hurdle in allowing that the group point of view is a function of the individual points of view (e.g., Levi, 1982).

Such was not Arrow's response to criticisms along the lines of Buchanan's. A "social preference" ordering does not represent the preferences of a *group agent*. On such a view, requiring social preference to be transitive is not to require rationality of a group, but of a rule to be put to certain uses. A rule for delivering social decisions, he says, "may be agreed upon for reasons of convenience and necessity without its outcomes being treated as evaluations by anyone in particular" (1951, p. xviii). That is, Arrow has in mind preferences without a preferring agent. Quoting Popper, he says that formulations of social preference can be "interpreted as hypostatizations of methodological rules" (1951, p. xix). Such "hypostatizations" can be used to articulate a notion of social welfare that could inform public officials in policy making.

4.2. Indeterminacy. Welfare is often identified with preference satisfaction in economics. The social choice framework presupposes that *social* welfare is a function of *individual* welfares. In my view, whether (a) only agents can be subject to rationality constraints and so social groups subject to Arrow's conditions must be agents, or (b) social welfare represents hypostatized methodology, we should appeal to *indeterminate preferences*. If (a) and the group agent has the goal of satisfying individual preferences, then conflict arises from the aims of satisfying conflicting individual preferences. If (b) and the methodological concern (in design or intended application) is with the preferences of in-group individuals, the same sort of conflict arises. That is, whether (a) or (b), if social choice aims to promote individual welfares, conflict in social preference results from conflict in individual preferences. It is conflict between various evaluations of the social options that occasions the need for a neutral point of view. How should we think of a *consensus* or neutrality among a set of preference orderings? The question is as relevant for agents of sufficient complexity as to endorse multiple values as it is for social choice. I explore the ramifications of dropping the assumption that rationality mandates resolution of (ordinal) conflict in preference before choice. Relaxing the uniqueness requirement allows us to identify *broader conceptions of consensus* than are available in the standard framework, as I will now explain.

A neutral or common ground or consensus position ought to preserve "shared agreements" among rival points of view and suspend judgment on controversial matters. This notion of consensus has been explored for full beliefs and probabilities as well as preferences (e.g., Levi, 1985; Seidenfeld et al., 1989; Elkin and Wheeler, 2016). The formal approach to preference aggregation that I propose below is motivated by this unified (across attitudes), normative conception of consensus and the resolution of disagreement. When there is a disagreement or conflict, instead of aggregating by majority rule or the like, parties to the disagreement may seek consensus, an initial common ground from which to engage in joint, non-question-begging inquiry or collective deliberation. The shared agreements constitute a neutral or non-controversial basis from which to pursue further inquiry that, in many cases, results in a resolution of the initial disagreement.

For beliefs and judgments of subjective probability, Levi marks a helpful distinction between consensus at the *outset* of inquiry—the initial shared agreements—and consensus as the *outcome* of inquiry—those beliefs and probability judgments jointly arrived at by experimentation, observation, etc. Indeterminacy is well-studied in the theory of subjective probability (e.g., Ellsberg, 1963; Levi, 1974; Gärdenfors and Sahlin, 1982; Walley, 1991; Seidenfeld et al., 2010). In the setting of expected utility, sets of probability functions have been explored as a representation of uncertainty that is not reducible to risk or to a unique probability distribution over the state space. Indeterminacy in probability, like indeterminacy in utility, can give rise to indeterminacy in values by generating multiple rankings of the options in terms of expected utility. The proposal below allows for a number of potentially conflicting values to be endorsed by an agent; standard social welfare functions do not. A prohibition against indeterminacy would be the demand that rational agents resolve any relevant such conflict by the time for choice.

Preferences, unlike qualitative beliefs but like subjective probability judgments, are not true or false (though there is a question of truth in *attributions* of certain preferences to agents). To suspend judgment between preference orderings, then, is not to suspend judgment concerning the truth value of the relevant orderings. Preference orderings represent ways of evaluating the feasible options. To suspend judgment between preference orderings, then, is to suspend judgment between ways of evaluating the options. Suspense may be thought of as regarding the rival preference orderings as *permissible* evaluations of the options. Equivalently, the relevant orderings are not *ruled out* in decision contexts. The intended interpretation of a set of preference orderings is that each element is *permissible* or *not ruled out* as an evaluation of the options.

Dewey and Tufts consider the case of a patriot and pacifist in wartime (1932, pp. 174-175). One set of the peaceful citizen's convictions require backing the state in the current war effort. On the other hand, his pacifist values commit him to opposition of the war. Until now, the story goes, the values of patriotism and pacifism had not been in conflict. Dewey and Tufts make clear that the issue is not one of weakness of will in which the citizen is tempted to act against what he judges to be best. He endorses both values and faces a genuine conflict. Examples of decision making with multiple goals abound. Consider a philosophy professor faced with a decision between three job offers. For professional aims, she judges department x to be better than department y

and department y better than department z. As far as considerations of her private life go, she orders the departments y over z over x. Resolving the conflict in her values revealed by her current dilemma may require developing value commitments that she currently lacks concerning how to weight professional and personal goals. That is to say, a definite way of trading off values induces an ordering of the options and is tantamount to a resolution in conflict.[4]

Dewey favors an assimilation of practical and theoretical reason. Nothing in the patriotic pacifist's values determines an all-things-considered best option. Such conflicts call for a form of *moral inquiry*, as he calls it, or *value inquiry* we might say more generally. Just as evidence can fail to decisively settle inquiry in favor of a particular answer, practical and moral considerations may fail to settle conflicts between goals and values. And just as in scientific inquiry, we ought not jump to conclusions or settle on a solution to the problem under consideration arbitrarily. That rational agents can suffer conflicts in value is not, it seems to me, controversial. But those tolerant of indeterminacy deny that rationality mandates all resolution of conflict by the time of choice (or that choice reveals such a resolution), just as they deny that, in inquiry, rationality mandates opinionation before the relevant evidence is in. According to the received view, rational agents are obliged to choose for the best, all things considered. Choosing for the best, all things considered, is not possible when there is no all things considered best option according to the agent's values. Arrow thinks that the ability to make consistent decisions is "one of the symptoms of an integrated personality," and on that basis, claims that the same ability is a symptom of an "integrated society" or group. But if "consistency" does not require a unique preference ordering for an individual, neither, the analogy would seem to suggest, is such a preference ordering required for a group (even if we resist attributions of agency to groups, but understand group preference as hypostatized methodology).

[4]Keeney and Raiffa begin their study of decision making with "multiple objectives" with a number of motivating examples (1993, pp. 1-5).

We will consider two concrete proposals for identifying general consensus positions with *sets* of preference orderings in what follows. One way is to consider what the shared agreements are concerning which ways of evaluating the options (preference orderings) are ruled out by the relevant parties. In that case, $\{R_i : i \in N\}$ is undisputedly a representation of group preferences; moreover, $\{R_i : i \in N\}$ excludes exactly the preference orderings that are commonly ruled out and suspends judgment on those in dispute. When the set is not a singleton, it is *indeterminate* what the social good is, or there is a conflict in social values deriving from a conflict in individual values. "Politics, it is usually agreed, is concerned with the common good," as Elster puts it (1997, p. 4). At the outset, though, the common good typically resists characterization in terms of a unique ordering (except for the very special case when the Pareto relation is complete), but is representable by a set of such orderings. A second set-based notion of consensus is considered in some detail below. Both proposals are intended to show the promise of indeterminacy in theorizing about social choice. That is, the particular examples I study here constitute a *proof of concept*.

Except when there is complete unanimity (the R_i are identical), $F(R_1, ..., R_n)$ represents neither consensus at the *outset* of inquiry and deliberation, nor consensus as the *outcome* of inquiry and deliberation. The output of $F(R_1, ..., R_n)$ is not simply the shared agreement among the n orderings. Similarly, the output of F does not represent the resolution of disagreement by deliberation or some form of inquiry, whether public or private depending on the individual or social application of the aggregation framework. Instead, $F(R_1, ..., R_n)$ is a form of political compromise. I do not mean to haggle over the term "consensus." One may insist, for instance, that an individual can consider herself a member of some group or party and *consent* to be represented by some value scheme that is established by means of some determinate aggregation method even if she does not personally endorse the scheme. But such an aggregation method typically rules out her way of evaluating the options, and so does not produce a consensus as shared agreement. It does not suspend judgment between the various candidate ways of evaluating the options. Some in the deliberative democracy camp, like Habermas and Rawls, have also found it fitting to distinguish this sort of "consensus" from consensus as shared agreement. The main point is that the conception under consideration in this essay differs substantially from alternative conceptions in important epistemological, decision-theoretic, and procedural ways that have consequences for collective choice.

Let me briefly summarize the reasons I see for exploring indeterminacy in the context of social choice. Foremost, in my view, is *neutrality*. If the objective of a social choice procedure is to satisfy the preferences of a certain group of individuals (whether the group should be thought of as an agent of not), satisfying the preferences of one individual can conflict with satisfying those of another, just as pursuing one goal can conflict with the pursuit of another for an individual agent. One notion of neutrality is given by suspending judgment among the rival ways of evaluating the options, not ruling them out for use in social or group choice. Such a notion is an analogue of seeking common ground in beliefs by suspending judgment on whatever is controversial while retaining the shared beliefs. Importantly, a neutral position does not beg questions, neither about what is true (in the case of belief) nor about which value assessments are permissible. Of equal importance, neutral positions understood in this weak way are not *arbitrary*. In the face of conflict, no recommendation is made to simply pick one belief, preference ordering, or probability function from among the rival candidates under consideration. Instead, the recommended procedure is to seek neutral ground from which to engage in further deliberation and inquiry.

The remainder of the chapter is given over to making a case for indeterminacy in social choice from a different angle. It will be shown that relaxing determinacy allows us to satisfy more of the central criteria for aggregation than the standard theory allows. While no proof can be offered for a particular normative view of social choice, results about aggregation can contribute to a case for or against certain aggregation formats and procedures. So, *if* certain axiomatic constraints on aggregation are found appealing, *then* the possibility results to come might be construed as considerations in favor of relaxing determinacy. Relatedly, allowing for indeterminacy allows us to resolve some important conceptual tensions in both democratic theory and welfare economics. These tensions are reflected in the impossibility of jointly satisfying certain constraints in the standard theory of social choice, and are, in some measure, resolved by relaxing determinacy.

5. Set-Valued Social Welfare Functions

"[T]he weakest link in Arrow's framework," Suzumura writes, "is his requirement of collective rationality" in the form of a single social preference ordering (1983, p. 100). From the beginning, Arrow recognized that there is no necessity in the assumption that social preferences are representable as a unique weak order of the options.

> Many writers have felt that the assumption of rationality, in the sense of a one-dimensional ordering of all possible alternatives, is absolutely necessary for economic theorizing [...] There seems to be no logical necessity for this viewpoint; we could just as well build up our economic theory on other assumptions as to the structure of choice functions if the facts seemed to call for it. (Arrow, 1951, p. 21)

Description and explanation are not the sole functions of a theory of collective choice. It could be, as I am urging in this chapter, that *rationality* does not require the assumption of a single weak ordering of all the alternatives. Even at the level of the individual, the assumption of a single preference ordering of the options is overly restrictive. In this essay, I will study *set-valued* social welfare functions, a generalization of element-valued social welfare functions allowing for indeterminacy:

$$\mathcal{F}\colon D \to \mathscr{P}(\mathfrak{R})$$[5]

Like Buchanan and Tullock, we are concerned to give consensus or shared agreement pride of place as a defensible basis from which to judge social good or at least preference (1962). We, however, adopt a more general conception of consensus. And like Sen, we consider a different aggregation format (1977). Social welfare functions are generalized to be set-valued. Attention will be focused on two concrete constructions of set-valued social welfare functions. There may well be other constructions worth considering (in fact, two additional ones are considered in the appendix), but

[5] $\mathscr{P}(\mathfrak{R})$ denotes the *powerset* of $\mathfrak{R}$, the set of all subsets of $\mathfrak{R}$.

these are two natural candidates. They are presented as a proof of concept for set-valued social welfare functions and indeterminacy in the context of social choice.

5.1. The Individual Orderings. The first is given very simply, if not seemingly very trivially, by the set of the n individual orderings themselves:

$$\mathcal{F}(R_1, ..., R_n) = \{R_i : i \in N\}$$

Such an $\mathcal{F}$ rules out none of the candidate ways of ordering the options. When there is no ordinal conflict among the individual preferences, $\mathcal{F}(R_1, ..., R_n)$ is a singleton because the candidate preference orderings coincide. Complete unanimity cannot be expected in general though.

One point to note is that each R_i trivially includes the weak Pareto relation, the agreements concerning strict preference: $\forall i \in N, \bigcap_{i \in N} P_i \subseteq R_i$. Actually, each R_i is an *extension* (sometimes called a *compatible extension* (e.g., Duggan, 1999)) of $\bigcap_{i \in N} P_i$. Extensions are formed by adding ordered pairs to the relation while preserving the asymmetric part (in the current setting, strict preference). The universal relation, $X \times X$, trivially completes any binary relation on X but does not generally preserve the asymmetric part of the relation. So the notion of a compatible extension is a more interesting notion of "extension" than mere subset inclusion for applications in decision theory, social choice, and related areas.

More formally, if R and R' are binary relations on X, R' is a (compatible) *extension* of R if (*i*) $R \subseteq R'$ and (*ii*) $R \backslash R^{-1} \subseteq R' \backslash R'^{-1}$. The notion of extension will figure prominently in the next constructions where we consider extensions of categorical preference.

5.2. Extensions of Categorical Preference. A given set of preference orderings may share certain preferences. For a profile of preference relations, $(R_1, ..., R_n)$, categorical strict preference is given by $\bigcap_{i \in N} P_i$; categorical indifference by $\bigcap_{i \in N} I_i$. We will call the union of these sets, $\mathcal{R}$, *categorical preference.*

$$\mathcal{R} = \bigcap_{i \in N} P_i \cup \bigcap_{i \in N} I_i$$

$\mathcal{R}$ represents the shared agreements or consensus among the R_i concerning both strict preference and indifference (Cf. Sen's IPAIR construction 2004b, p. 674). Certain options may be non-comparable under the categorical preference relation, i.e., x may not be categorically preferred, inferior, or indifferent to some option y. In general, categorical preference is not the same as taking the intersection of the individual orderings. The following example serves to illustrate this point.

EXAMPLE 1. *Consider two preference relations, R_1 and R_2, and let $X = \{x, y, z\}$. Let R_1 be such that xP_1yP_1z, and R_2 be such that xP_2yI_2z. We have $(y, z) \in R_1$ and $(y, z) \in R_2$, but $(y, z) \notin \bigcap_{i=1,2} P_i$ and $(y, z) \notin \bigcap_{i=1,2} I_i$. That is, $(y, z) \notin \mathcal{R}$ yet $(y, z) \in \bigcap_{i=1,2} R_i$.*

Taking the intersection of the individual orderings can deliver universal strict preference for one option over another despite some individual orderings counting the two options as indifferent. Some authors see this as a desirable feature. In Example 1, the social choice of y instead of z makes no individual worse off and makes one individual better off. But our concern at present is *categorical preference*, the precise set of unanimous preferences (but see the appendix for an exposition of Levi's view that categorical preference ought to preserve categorical weak preference in addition to categorical strict preference and categorical indifference). Categorical preference as identified here is a subset of the intersection of the individual orderings.

PROPOSITION 21. *For all $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\mathcal{R} \subseteq \bigcap_{i \in N} R_i$.*

Like $\bigcap_{i \in N} R_i$, $\mathcal{R}$ is not generally complete. If there exists $i, j \in N$ such that xP_iy and yP_jx, for instance, $\mathcal{R}$ will include neither (x, y) nor (y, x). Categorical preference, then, is not a universally applicable criterion of *admissibility* or permissible choice. The same is true of the Pareto relation (one formulation of which is our categorical strict preference $\bigcap_{i \in N} P_i$). It is in general incomplete and consequently silent on many menus of options. If the function of preference is to determine acceptable options for choice, it is important to supplement incomplete preferences. A standard move is to consider some Pareto-extension rule that completes the relation.

Pareto-Extension Rule. For all $x, y \in X$, xRy *iff* $(y, x) \notin \bigcap_{i \in N} P_i$.

Pareto-extension rules are essentially arbitrary, singling out a particular extension of the Pareto relation (not a weak ordering) when there are many possible extensions. On what basis are the remaining possibilities ruled out? The set of weak order extensions represents the potential compromises between the parties to the aggregation. By also including all individual orderings, the set of possible extensions represents well the social diversity in preferences. (Sen records a number of limitative results for the Pareto-Extension Rule (1970a, pp. 74-76).)

A second set-based notion of *consensus* is given by the (typically distinct) set of weak order extensions of $\mathcal{R}$. The motivation behind considering this set is that, in suspending judgment between candidate ways of evaluating the options, none of the potential ways of resolving the conflict between those candidates should be ruled out.

Various results show that certain binary relations can be extended to certain sorts of orderings (e.g., Szpilrajn, 1930; Hansson, 1968; Fishburn, 1973; Suzumura, 1983; Duggan, 1999). We will appeal to the following result for quasiorders, i.e., transitive and reflexive binary relations.

THEOREM 19. *(Hansson, 1968, pp. 454-455) Every quasiorder has a weak order extension.*

Categorical preference as defined here is a quasiorder.

PROPOSITION 22. *For all* $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\mathcal{R} = \bigcap_{i \in N} P_i \cup \bigcap_{i \in N} I_i$ *is a quasiorder.*

So Hansson's extension theorem applies.

PROPOSITION 23. *For all* $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\mathcal{R}$ *has a weak order extension.*

As Pedersen, for example, emphasizes, weak order extensions of quasiorders are not in general unique (Pedersen, 2009b, p. 11). If for some $x, y \in X$ and some quasiorder R, neither xRy nor yRx, then there is a weak order extension R' of R such that $xR'y$ and another extension R'' of R such that $yR''x$. So the set of weak order extensions of a quasiorder is not generally a singleton.

Let $Ext(\mathcal{R})$ be the set of weak order extensions of $\mathcal{R} = \bigcap_{i \in N} P_i \cup \bigcap_{i \in N} I_i$:

$$Ext(\mathcal{R}) = \{R \in \mathfrak{R} : (i)\ \mathcal{R} \subseteq R, \text{ and } (ii)\ \mathcal{R} \backslash \mathcal{R}^{-1} \subseteq R \backslash R^{-1}\}$$

$\mathcal{R}$ represents exact agreement on strict preference and indifference. The set of all weak order extensions of $\mathcal{R}$ represents ways of resolving disagreements in such a way as to preserve shared agreements. That is, $\mathcal{R}$ includes all the ways of extending exact agreement to a preference ordering. The second concrete construction that we will consider in this chapter is given by $Ext(\mathcal{R})$:

$$\mathcal{F}(R_1, ..., R_n) = Ext(\mathcal{R})$$

Observe that for each $i \in N$, $R_i \in Ext(\bigcap_{i \in N} R_i)$ (Proposition 24 below). Under sufficient divergence of preferences, it is possible that $\mathcal{R} = \Delta = \{(x, x) : x \in X\}$. In that case, $Ext(\mathcal{R}) = \mathfrak{R}$. That is, no social ordering of the alternatives can be ruled out since every ordering is a resolution of the (nearly total) disagreements that preserves the (nearly null) shared agreements.

Is there a principled case for focusing on $\mathcal{R}$ rather than $\bigcap_{i \in N} R_i$ or some Pareto relation such as $\bigcap_{i \in N} P_i$ as a basis for generating a consensus set of preferences? I think there is. $\mathcal{R}$ more precisely captures consensus than does $\bigcap_{i \in N} R_i$. As Example 1 demonstrates, categorical preference consists of all and only explicit agreements (in both strict preference and indifference), while $\bigcap_{i \in N} R_i$ can introduce strict preferences that are *not* unanimously shared. Whereas $\bigcap_{i \in N} R_i$ sees agreements that are not there, thereby overly constraining the set of weak order extensions, $\bigcap_{i \in N} P_i$ does not properly recognize the extent of agreement, thereby failing to sufficiently constrain the set of extensions. If common ground or shared agreement among some set of candidate ways of evaluating the options requires suspending judgment among the candidate orderings by not ruling them out, then, as the following proposition attests, $\bigcap_{i \in N} R_i$ cannot be taken to generate consensus as common ground or shared agreement.

Proposition 24. *For all* $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\{R_i : i \in N\} \subseteq Ext(\mathcal{R})$, *but it is* ***not*** *the case that for all* $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\{R_i : i \in N\} \subseteq Ext(\bigcap_{i \in N} R_i)$.

If there were a reason to think that only agreement in *strict* preferences mattered, that no part of consensus mandates preserving agreement about *indifference* between options, then $Ext(\bigcap_{i \in N} P_i)$ would be of real and living concern as an account of consensus or neutrality in preference. I know of no such reason. Consequently, elements of $Ext(\bigcap_{i \in N} P_i) \backslash Ext(\mathcal{R})$ may fail to represent legitimate ways of resolving conflicts in values since they fail to preserve some agreements (but see the appendix for some indications of how an account in terms of categorical strict preference could be worked out). $\{R_i : i \in N\}$ and $Ext(\mathcal{R})$ are candidates for lower and upper bounds, respectively, for set set-valued approaches to ordinal preference aggregation that are motivated by neutrality.[6] My purpose, however, is just to indicate some plausible models that can be used in the consistency proofs in the following section.

6. Some Possibility Results

The above conception of consensus in preferences can be partly substantiated by *possibility* results. My hope is that these results provide a new philosophical perspective on aggregation and a reason to consider indeterminacy in social preference. Deliberative democrats should, it seems to me, receive this as good news, as I will explain in Section 7.

[6] In a sense, $Ext(\mathcal{R})$ is the preference analogue of the convex hull of a set of probability functions. $Ext(\mathcal{R})$ includes more orderings than $\{R_i : i \in N\}$ just as the $conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ includes more probability functions that $\{\boldsymbol{p}_i : i \in N\}$. And in analogy to interpretation sometimes offered of the convex combinations of the $\boldsymbol{p}_i$, $i \in N$ (Levi, 1980, 2009), the $R \in Ext(\mathcal{R})$ are presented as potential resolutions of the conflict among R_i, $i \in N$.

For set-valued social welfare functions, the two concrete recipes for aggregation considered above serve as examples of constructions satisfying (extensions of) both the Arrow conditions and the conditions involved in Sen's impossibility result. First, the axioms for social welfare functions need to be suitably extended to accommodate indeterminacy. While these extensions are genuine generalizations, reducing to the standard versions of the axioms when $\mathcal{F}(R_1, ..., R_n)$ is a singleton, there are typically multiple ways to generalize an axiom. In each case, I submit that the essence of the axiom, whatever is compelling about it, has been preserved, but readers will need to make similar judgments for themselves.

U does not need to be altered because we may still take the domain of $\mathcal{F}$ to be $\mathfrak{R}^n$.

(**U**) The domain of F is the set of all possible profiles, i.e., $D = \mathfrak{R}^n$.

Since $\mathcal{F}$ takes its range in *sets* of orderings, however, for Pareto we have the following:

(**P**) For all $x, y \in X$ and all profiles in D, if xP_iy for all $i \in N$, then xPy for all $R \in \mathcal{F}(R_1, ..., R_n)$.

Independence of irrelevant alternatives needs to be similarly generalized.

(**IIA**) For all $x, y \in X$ and any $(R_1, ..., R_n), (R'_1, ..., R'_n) \in D$, if xR_iy iff xR'_iy for all $i \in N$, then xRy for some $R \in \mathcal{F}(R_1, ..., R_n)$ iff $xR'y$ for some $R' \in \mathcal{F}(R'_1, ..., R'_n)$.

(**IIA**) could be strengthened by swapping out the existential quantifiers in the consequent for universal quantifiers: xRy for *all* $R \in \mathcal{F}(R_1, ..., R_n)$ iff $xR'y$ for *all* $R' \in \mathcal{F}(R'_1, ..., R'_n)$. The choice between these formulations of (**IIA**) will not make a difference for the results to come. For non-dictatorship, we put

(**D**) There is no $i \in N$ such that, for all $x, y \in X$ and all $(R_1, ..., R_n) \in D$, if xP_iy, then xPy for all $R \in \mathcal{F}(R_1, ..., R_n)$.

We are now in a position to state a possibility result for set-valued SWFs. It turns out that both of the concrete structures under consideration here satisfy the above generalizations.

Proposition 25. *Suppose $|X| > 2$. There are set-valued social welfare functions that satisfy* ***U***, ***P***, ***IIA***, *and* ***D***.

That $\mathcal{F}(R_1, ..., R_n) = \{R_i : i \in N\}$ satisfies the conditions—which is easy enough to check—suffices to establish the proposition. The proof, however, is for the slightly more interesting case $\mathcal{F}(R_1, ..., R_n) = Ext(\bigcap_{i \in N} R_i)$.

Sen's liberalism constraint itself is the subject of outstanding controversy, so some degree of controversy concerning the extension of the constraint to the set-valued context is unavoidable. Since $\mathcal{F}(R_1, ..., R_n)$ is not typically a singleton, we have to consider how to represent the constraints on social preference generated by Sen's liberalism property. One straightforward generalization is to require that *some* permissible social preference relation reflect an individual's strict preference over some pair of options.

(**L**) For all $i \in N$, there is a pair of options $x, y \in X$ such that xP_iy implies xPy and for some $R \in \mathcal{F}(R_1, ..., R_n)$, yP_ix implies yPx for some $R \in \mathcal{F}(R_1, ..., R_n)$.

Why not require that *all* $R \in \mathcal{F}(R_1, ..., R_n)$ respect i's strict preference for some pair of options? Because doing so is at odds with interpreting $\mathcal{F}(R_1, ..., R_n)$ as a neutral position for the R_i. When there is ordinal conflict, the approach explored in this chapter counsels suspending judgment between rival ways of evaluating the options, not automatically imposing one individual's preferences on another. Some democratic theorists argue that the relevant sorts of procedural constraints for democracy require that each individual have his or her interests given equal consideration in social choice (e.g., Dahl, 1989, p. 166), not that social preference automatically falls in line with certain individual preferences. **L** requires respect for individual interests by requiring that they *not be ruled out* at least for a pair of options. In fact, for examples of $\mathcal{F}$ under consideration, this is true for more than just two options. These aggregation procedures may be thought of as fair in the sense that they exclude none of the individual preferences in the profile from the "social" preferences.

Proposition 26. *There are set-valued SWFs that satisfy* ***U****,* ***P****, and* ***L****.*

If we take the conditions involved in Arrow's and Sen's results as desiderata or normative criteria for preference aggregation, Propositions 25 and 26 are points in favor of relaxing determinacy for aggregation.

7. Deliberative Democracy vs. Social Choice Theory

According to Dryzek and List, the "two most influential traditions of contemporary theorizing about democracy, social choice theory and deliberative democracy, are generally thought to be at loggerheads, in that the former demonstrates the impossibility, instability or meaninglessness of the rational collective outcomes sought by the latter" (2003). Recall just how tremendous the tension between social choice theory and deliberative democracy is alleged to be. Besides Hardin's claim that social choice theory reveals "flaws—grievous foundational flaws—in democratic thought and practice" (1993), we have the following assessment from William Riker, for example.

> The main thrust of Arrow's theorem and all the associated literature is that there is an unresolvable tension between logicality and fairness [...] no adequate resolution of the tension has been discovered, and it appears quite unlikely that any will be. The unavoidable inference is, therefore, that, so long as a society preserves democratic institutions, its members can expect that some of their social choices will be unordered or inconsistent. And when this is true, no meaningful choice can be made. (1982, p. 136)

Mackie cites many others drawing similar conclusions from Arrow-type results in social choice theory (Mackie, 2003).

The first point that I wish to make is that social choice theory does not seem nearly as destructive of ideals concerning democratic consensus *once we allow for more general notions of consensus* as the set-valued framework does. The impossibilities confronting the standard framework, the allegedly "unresolvable tension between logicality and fairness," give way to certain possibility results for sets of attractive aggregation conditions like those involved in Arrow's and Sen's results. Set-valued aggregation functions admit normative motivations in terms of *common ground* or *neutrality* when preferences conflict. So, the foregoing account of preference aggregation presents a possible way of reconciling the two main traditions of theorizing about democracy. The question now is, to what extent does indeterminacy allow us to accommodate deliberative ideals into the social choice framework?

Jürgen Habermas is one of the foremost champions of deliberative democracy. Elster characterizes his position: "the goal of politics should be rational agreement rather than compromise, and the decisive political act is that of engaging in public debate with a view to the emergence of consensus" (Elster, 1997, p. 3). Social decisions, according to Habermas, can be justified from an uncoerced consensus emerging through deliberation. Similarly, Rawls appeals to the consensus that would emerge in a hypothetical neutral position, the veil of ignorance, to justify policies, principles, and institutions.

According to the way Elster sees Habermas's views on deliberative democracy, "there would not be any need for an aggregating mechanism, since a rational discussion would tend to produce unanimous preferences" (Elster, 1997, p. 11). Like Elster and Dryzek and List, I suspect this to be overly optimistic, if for no other reason than there often is no time to carry the deliberation through to its conclusion, just as there may be no time to collect all of the evidence if there were a deadline to report findings in some theoretical inquiry. Dryzek and List conclude "so even after discussion-induced preference changes aggregation of conflicting preferences may be necessary" (2003, p. 6). Set-valued SWFs would seem the appropriate representation of group preference, not only at the outset of group deliberation and inquiry, but at any subsequent stage at which complete unanimity has not been achieved. What has been offered here is a way of "aggregating" potentially conflicting preferences that respects the conflicts and does not presume to resolve it in the absence of deliberation simply because a choice must be made. Neutrality among conflicting preferences does not rule out candidate preference orderings. In such cases, consensus can be represented by the set of permissible orderings, those that have not been ruled out. One way to summarize the results of Section 6 is that attractive aggregation constraints are jointly satisfiable for models of aggregation that are neutral in the indeterminate sense.

But where does deliberation fit in? I have proposed no models of deliberation. Instead, what I hope to have done is cleared ground for deliberation in the social choice setting. $\mathcal{F}(R_1, ..., R_n)$ does not represent as much common ground as it is possible to achieve because individuals might change their preferences through deliberation. Bohman writes, "For a democracy based on public deliberation presupposes that citizens or their representatives can take counsel together about what laws and policies they ought to pursue as a commonwealth. And this in turn means that the plurality of competing interests is not the last word, or sole perspective, in deciding matters of public importance. The problem, to use Kant's terms, is to bring about 'the public use of reason'" (1997, p. x). I will sketch an outline of how deliberative ideals might enter into the generalized social choice framework under discussion.

Consider an analogy with preceding chapters. In the theory of subjective probability, there are widely celebrated convergence and merging of opinion results (e.g., Savage, 1954; Blackwell and Dubins, 1962). These are results to the effect that, under just a few assumptions, two probability measures that may initially disagree quite substantially agree in the limit of updating on a shared (infinite) stream of data.[7] In the long run of inquiry, differences in priors "wash out." At any finite stage, some disagreement may persist. It is possible that given some inference or decision problem facing the group, there is no time or opportunity to collect more evidence. As shown in Chapters 2 and 3, in such cases, aggregation methods can be employed to arrive at a group consensus that is a functional, non-question begging point of view from which to make inferences and decisions.

I want to suggest an analogous outlook on deliberation. Structural similarity between deliberation and inquiry should be expected by those sharing Dewey's view of the unity of practical and theoretical reason. We can distinguish between consensus at the *outset* of deliberation and

[7] Recall that in the case of Blackwell and Dubins' result, those assumptions include that the probability measures are *mutually absolutely continuous*, i.e., they assign probability 0 to the same events, and that the conditional probabilities for each are *predictive*, i.e., each measure admits a *regular conditional distribution* for each finite history.

consensus as the *outcome* of deliberation just as we did for inquiry. At the outset of deliberation, $\mathcal{F}(R_1, ..., R_n)$ represents common ground. If deliberative democrats are right that deliberation tends to produce agreement, then the size of $\mathcal{F}(R_1, ..., R_n)$ should decrease when the R_i are subjected to deliberation. At any point—importantly, even in the absence of unanimity—indeterminacy allows us to identify a general consensus, one that suspends judgment among candidate ways of evaluating the options. A central feature of deliberative democracy, according to Larmore, is "the pursuit of a common ground on which people can stand despite their deep ethical and religious differences" (2003). At least for the sake of the argument[8], sets of preference orderings can be identified which can serve important functions as common ground, even while individuals retain their potentially diverse values. Moreover, as we have seen, conceptions of consensus along indeterminate lines have other attractive features that determinate ones lack. But supplementing my proposal for indeterminate preference aggregation with an account consensus-strengthening deliberation (and lexicographic choice procedures (Section 9.2)) may serve to help forestall worries about "democratic paralysis" and the permissive nature of choice under indeterminacy.

I confess that I do not yet have an account to offer of how the analogue process for deliberation would go.[9] But some form of substantive exchange among the individuals in the group has long been recognized as a supplement or perhaps an alternative to social choice theory even by those outside of the deliberative democracy camp. Sen, for example, writes that

> it can certainly be argued that the eventual guarantee for individual freedom cannot be found in mechanisms of collective choice, but in developing values and preferences that respect each other's privacy and personal choices. (1970a, p. 85)

Elster associates this idea with deliberative democracy.

> The input to the social choice mechanism would then not be the raw, quite possibly selfish or irrational preferences that operate in the market, but informed and other-regarding preferences (1997, p. 11).

According to views like those expressed here in Sen's and Elster's remarks, elsewhere by Dryzek and List (Dryzek and List, 2003), for example, deliberation or some form of exchange *shapes* individual

[8] An alternative account would be to require that individuals genuinely adopt $\mathcal{F}(R_1, ..., R_n)$ as their system of values.

[9] I do not mean to deny that a lot of relevant work by others on the topic exists.

preference. So, eventually, the aggregation function's input will be limited to certain sorts of profiles, perhaps even relatively homogenous ones. But supposing that a clear and compelling account of how deliberation reduces disagreement or restricts the domain could be articulated does not obviate the need for methods of formulating neutrality or consensus at points when unanimity in the sense of a single preference ordering has yet to be achieved. At any point, parties may have conflicting preferences. (Others have suggested reconciliations between deliberative democracy and social choice that appeal to both deliberation and post-deliberation aggregation, though they differ in substantial ways from the proposals made here (see, e.g., Dryzek and List, 2003; Perote-Peña and Piggins, 2015).) In fact, even the merging of opinions results for probability secure only similarity of opinions in any finite time, not complete unanimity.

According to deliberative democrats, consensus is the politically legitimate basis for collective decision making. In solidarity with Habermas's and Rawls's visions, the conception of consensus studied here has it that on rational (as opposed to descriptive or explanatory) models of preference transformation, mere aggregation is no substitute for deliberation for resolving conflict. The set $\{R_i : i \in N\}$, for example, may be taken to represent consensus as shared agreement in preferences at the outset of deliberation. Securing initial consensus matters for establishing a common platform of values from which to pursue policy decisions, resolutions in value conflict, and social choice generally. Consensus at the outcome of deliberation may yet resist representation by a unique preference ordering. But through deliberation, inquiry, public reasoning, and the like, a stronger consensus position will ideally emerge.

8. Welfarism vs. Non-Welfarism

Welfarism is the view according to which social policies and outcomes ought to be judged according to the consequences that they have for individuals. Widening the informational basis by including more than merely ordinal information about individual preference is a sophistication of the social choice framework along welfarist lines. Non-welfarist approaches are concerned with consequences that are not restricted to individual preferences, welfare, or utility. Social choice should, for example, take into account *procedural* information. According to Suzumura, Sen's result exposes the conflict between welfarism in the form of the Pareto condition, **P**, and non-welfarism in the form of social respect for individual libertarian rights, **L** (Bossert and Fleurbaey, 2015, pp. 186-188). No social decision rule satisfies both **P** and **L** given Universal Domain, **U**.

Allowing for indeterminacy in the social welfare function resolves some measure of the tension between the two sorts of considerations. To the extent that the tension is resolved by allowing for indeterminacy, the conflict between welfarism and non-welfarism is (at least in part) an artifact of the assumption of determinacy in social preference. As we saw above, **P** and **L** *are* jointly satisfiable (in the presence of **U**) in the framework of indeterminate social choice, unlike in the standard theory.

Since the tension between welfarism and non-welfarism may rest on Sen's result to a lesser degree than the tension between deliberative democracy and social choice theory rests on Arrow's result, the claim of resolution may need to be more modest. Recall, moreover, that there may be reservations concerning the extension of **L** to settings allowing for indeterminacy. There may be reservations about whether our extension of **L** continues to represent an attractive procedural, non-welfarist constraint—at least when $\mathcal{F}(R_1, ..., R_n)$ is not a singleton; when it is a singleton, our extension of **L** collapses to Sen's original formulation. A resolution of this issue will depend on whether or not we accept an account that has it that what is fair is for $\mathcal{F}$ not to rule out any individual ordering absent due process, instead of an account according to which certain individual preferences must be automatically honored. Due process here would be some form of value inquiry, undertaken without begging questions or prejudging the issue of permissible orderings.

9. Some Challenges

9.1. Limitative Results. The name of the game in social choice theory is impossibility, and there have been many bright players. But it is clear that not *every* impossibility result should alarm us. Such results are disconcerting only to the extent that the joint satisfaction of the relevant constraints is judged compelling. We now turn to some limitative results for set-valued SWFs for which questions of such judgment are relevant.

One approach to coping with Arrow's result is to weaken the assumption of collective rationality, i.e., ordering. A limitation frustrating many such attempts is the existence of a vetoer. Individual $i \in N$ is called a *vetoer* for F if for all $x, y \in X$ and all $(R_1, ..., R_n) \in D$, xP_iy implies xRy (where $R = F(R_1, ..., R_n)$). In words, i has the veto if i can always block any option from being strictly preferred. Another limitation found worrisome is the existence of an *oligarchy*, a sort of generalization of the notion of dictatorship to a group. F is called *oligarchic* if there exists a set $M \subseteq N$ such that for all $(R_1, ..., R_n) \in D$ and all $x, y \in X$, (i) if xP_iy for all $i \in M$, then xPy, and (ii) if xP_iy for some $i \in M$, then xRy (where $R = F(R_1, ..., R_n)$ and $P = R \backslash R^{-1}$). In words, a group is an oligarchy if it is decisive for F relative to any profile, and each member of the group is a vetoer.

Again, there are various ways to lift these definitions to the context of set-valued aggregation. What is more, certain of these generalizations clearly apply to the constructions studied in the preceding sections. This should be neither surprising nor disconcerting given the philosophical motivations that underwrite set-value aggregation as I have presented it. Consider, for example, the following generalization of the notion of a vetoer. Individual $i \in N$ is called a *vetoer* for $\mathcal{F}$ if for all $x, y \in X$ and all $(R_1, ..., R_n) \in D$, xP_iy implies xRy for some $R \in \mathcal{F}(R_1, ..., R_n)$.[10] On the two constructions considered above, *every* individual has the veto. If instead we require that xRy for *all* $R \in \mathcal{F}(R_1, ..., R_n)$ in the consequent of the definition of a vetoer, no individual would be a vetoer. And this (in conjunction with the possibility results) shows that no collection of the axioms under consideration imply the existence of *that sort* of vetoer in the context of set-valued aggregation.

[10]Notice that, in the setting of choice functions considered below and in the Appendix, for a choice function generated from $\mathcal{F}(R_1, ..., R_n)$ by V-admissibility, vetoers insure the admissibility of options they strictly prefer.

The notion of oligarchic aggregation can be similarly generalized. For example, call $\mathcal{F}$ *oligarchic* if there exists a set $M \subseteq N$ such that for all $(R_1, ..., R_n) \in D$ and all $x, y \in X$, (i) if xP_iy for all $i \in M$, then $(x, y) \in R \backslash R^{-1}$ for all $R \in \mathcal{F}(R_1, ..., R_n)$, and (ii) if xP_iy for some $i \in M$, then xRy for some $R \in \mathcal{F}(R_1, ..., R_n)$. This definition also applies to the concrete constructions considered above. The whole set N is the oligarchy: $M = N$. But this limiting case robs the designation of its sting. If the definition of oligarchic aggregation is altered so as to change the notion of vetoer in clause (ii) as we did just above, neither aggregation model presented is oligarchic in this latter sense. Again, this shows that no collection of the axioms under consideration imply that the set-valued aggregation function is oligarchic in the latter sense since we have counterexamples.

But what about the first, weaker generalization of the notion of a vetoer and the corresponding generalization of oligarchic aggregation? If aggregation is interpreted as providing *consensus* or *common ground*, the existence of vetoers should be met with equanimity. Insofar as the vetoer's preferences are not ruled out, they are regarded as permissible. The neutral or consensus position represents suspension of judgment among the permissible evaluations of the options by not excluding them.

9.2. Choice under Indeterminacy. A consensus, however weak, is always identifiable in the setting of indeterminate preferences. But a concern that immediately presents itself is how to use a set of orderings for the purposes of decision making. At first sight, the problem of social choice is exacerbated by allowing for indeterminacy. Numerous, perhaps the majority, of social decisions require choice before we exhaust the means for resolving conflicts in values, to say nothing of potentially interminable conflict. How are choices to be made from the potentially numerous options licensed by indeterminate preferences? Isn't indeterminacy concession to "democratic paralysis"?

In cases of indeterminacy in preference, two things are called for.[11] First, we need a decision theory that applies to cases of unresolved conflicts in values, not the pretense or legislation that all such conflicts have been or are to be resolved prior to choice. If an individual agent with multiple goals faces a decision problem for which those goals are relevant and conflicting, rationality does not require the immediate resolution of conflict or determinacy of preference. The same goes for groups. As Levi puts it, "the need to take decisions (which, in my view, is as urgent in pure research as it is in practical deliberation) does not mandate or even excuse unjustified resolution of conflict or leaping to conclusions" (1982). The second thing needed is an account of the rational resolution of conflicts in value, some generalization of Dewey's moral inquiry. I took up the second in Section 7. I consider the first issue in the present section.

In many cases, a number of options on a menu are *admissible* or acceptable for choice relative to current value commitments. Let $\mathcal{X}$ be a collection of non-empty subsets of non-empty X. $\mathcal{X}$ is the set of *menus*. A *choice function* is a map $C : \mathcal{X} \to \mathscr{P}(X)\backslash\{\varnothing\}$ such that $C(Y) \subseteq Y$ for all $Y \in \mathcal{X}$. Often, studies of choice functions focus on the set of all *finite*, non-empty subsets of X. For simplicity, I will also. One criterion of decision for indeterminate preferences restricts choice to options that are optimal with respect to some permissible ordering of the options, called *value admissibility* or *V-admissibility* for short (Cf. Levi, 1986a).

V-Admissibility. An option $x \in X$ is *V-admissible* in $Y \in \mathcal{X}$ relative to a set of orderings $\mathcal{O}$ if, for some $R \in \mathcal{O}$, xRy for all $y \in Y$.

The set of V-admissible options are those options on the menu that are best according to *some* permissible preference ordering, i.e., some way of evaluating the options that has not been ruled

[11] The topic of incomparability and conflict is of broad interest in value theory. Ruth Chang, for example, has written a lot about choices between options that are *on a par*. According to Chang, options are on a par when they are not comparable by strict preference or indifference but by *some other* relation. (In the case of value conflict, the options *are comparable* under each permissible ordering, even if they are not comparable under any *categorical* preference relation.) Her suggestion is to create new reasons for favoring one option over another in order to break the parity (e.g., 2009). Sven Ove Hansson considers a *development account* according to which preferences can be genuinely incomplete and new preferences must be *developed* to cover incomparable options (1993). In much the same spirit as the developments in this chapter, Sen says, "A particular reason for incompleteness may be the presence of *multiple preferences* between which a person cannot decide" (2004b, p. 669).

out. There are other choice rules for sets of preference orderings. For instance, an option $x \in \mathcal{X}$ is *maximal* for $\mathcal{O}$ if for all $y \in Y$ there is some $R \in \mathcal{O}$ such that xRy. Unlike V-admissibility, there need not exist one $R \in \mathcal{O}$ such that xRy for all $y \in Y$ for x to be maximal. It is easy to see that maximality licenses choices that are not optimal according to *any* permissible preference ordering in $\{R_i : i \in N\}$. I will restrict my attention to V-admissibility in this chapter. V-admissibility would seem the appropriate standard for a set of orderings each of which is regarded as *permissible* in assessing optimality. A set of preference orderings determines a choice function *via* V-admissibility. Let $\mathcal{O}$ be a set of preference orderings:

$$C_{\mathcal{O}}(Y) = \{x \in Y : \exists R \in \mathcal{O}\ \forall y \in Y\ xRy\}$$

Often, even for single agents and even when $\mathcal{O}$ is a singleton, $C_{\mathcal{O}}(Y)$ will not be a singleton. In the presence of indeterminacy—that is, when $\mathcal{O}$ is not a singleton—the agent cannot be assumed to be *indifferent* between the options in the choice set. Yet, there are circumstances that force the selection of a single option anyway. The philosophy professor choosing between three jobs may not have the opportunity to resolve her conflict between personal and professional goals. Dewey and Tufts' citizen may enlist in the war effort while recognizing that his choice is not for the best according to his pacifist commitments. In each case, the choice is optimal relative to some permissible way of evaluating the options, but the conflict in values has not been resolved.

9.2.1. *Lexicographic Preferences.* If a set of conflicting values completely describe an agent's value commitments, then, without some account of conflict resolution, the story ends with a set of admissible options, all of which are optimal relative to some permissible way of evaluating the options but none of which is best, all things considered. But the story often does not end there. Besides undertaking further efforts to resolve conflicts in value (which I argued above is a role for deliberation in social interpretations of aggregation), agents may have values described by *lexicographic* preference structures. That is, an agent may employ *higher-order* values to help break ties. By "tie-breaking" I do not mean picking between options concerning which the agent is indifferent. The tie is in admissibility, not weak preference. For example, Dewey and Tufts' character might employ a further value such as honoring his mother's wishes on the matter. His commitments may prevent him from allowing such a criterion to figure into the determination of admissibility at the first stage.[12] Similarly, groups may employ second (third, etc.) tier criteria to winnow the set of admissible options. In some cases, the higher-order values may make little difference, but in some cases they will. For example, these criteria may include types of voting in certain circumstances. There may be multiple permissible orderings at any tier. Applying V-admissibility for permissible orderings at higher tiers, (potentially) further restricts the set of options deemed admissible at lower tiers. Just as with first-tier values, I do not assume that it is *mandatory* to hold a particular higher-order value. The existence of higher-order values and criteria for choice is a matter of what values the agent, group, or methodology is committed to.

I now want to address two worries about the account of choice just sketched. First, is the account reducible to revealed preference or some other binary comparisons after all? Second, does the account lead to violations of appealing standards of rational choice? In connection with the second question, I will consider the injunctions against choosing dominated options and susceptibility to money pumps.

9.2.2. *Choice, Preference, and Voting.* Under indeterminacy, choice does *not* reveal preference.[13] V-admissibility is not reducible to binary comparisons of the options. So, admissibility cannot be represented by a single binary relation. There are choice functions defined from a set of orderings *via* V-admissibility that are not representable as choosing the optimal or maximal

[12] In the setting of expected utility, when different permissible probability distributions rank different options as optimal, a popular second-tier criterion is *security*. The options that are (first-tier) admissible in expectation (E-admissible) are those that maximize expectation relative to *some* permissible probability distribution. Rank the options that are admissible in expectation according to their lowest expectation across permissible probability distributions. Select from the E-admissible options those with the best worst-case outcome (e.g., Ellsberg, 1963; Levi, 1986a).

[13] There is a generalization of rational choice theory, Helzner's *conditional choice*, that appeals to *sets* of relations to rationalize a (conditional) choice function (Helzner, 2013). Given the ability to elicit sufficiently rich *conditional* judgments of admissibility—judgments relative to hypothetical admissibility-determining informational and valuational states—the relationship between (conditional) choice and revealed preference requires reassesment.

elements according to some binary relation R (see also Example 3 below and Example 4 in the appendix). A binary relation R on X is said to *rationalize* a choice function if, for all $Y \in \mathcal{X}$,

$$C(Y) = \{x \in Y : xRy \text{ for all } y \in Y\}.$$

EXAMPLE 2. *Let $X = \{x, y, z\}$ and $\mathcal{O} = \{R_1, R_2\}$ with xP_1yP_1z and zP_2yP_2x. Using V-admissibility, we have $C(\{x, y\}) = \{x, y\}$, $C(\{y, z\}) = \{y, z\}$, $C(\{x, z\}) = \{x, z\}$, and $C(\{x, y, z\}) = \{x, z\}$. Suppose that there is an R that rationalizes C. Then, yRx and yRz (from binary choices). Yet, $y \notin C(\{x, y, z\})$, contrary to the assumption that R rationalizes C. It follows that no binary revealed preference relation rationalizes C.*

In the appendix, this point is made more general and precise with the help of some auxiliary concepts from revealed preference theory and the framework of choice functions. For the sake of readability and ease of exposition, however, I attempt to confine what formal material about choice functions I can to the appendix. The important point for us is that the account of choice advocated here departs from revealed preference.

Those of us tolerant of indeterminacy are not alone in voicing criticisms of revealed preference theory. Though behaviorism has fallen out of fashion in psychology, it continues to play a prominent role in economics. Sen (1973) traces some of the history of the behaviorist origins and motivations for revealed preference theory. But by now, many have come to reject or be quite skeptical of arguments for revealed preference theory that rest on behaviorist premises (e.g., Sen, 1973; Hausman, 1992; Mongin, 1997; Dietrich and List, 2016). According to Mongin, the simple point that what one chooses may be distinct from what one prefers "had to be made repeatedly by philosophers against those economists who defend 'revealed preference theory' and thus confuse decision, and indeed only one special kind of decision—i.e., a choice between alternatives—with preference" (1997, pp. 3-4). Sen makes a related point. In certain circumstances, such as the much-discussed prisoner's dilemma, the "behaviour pattern that will make each better off in terms of their real preferences is not at all the behaviour pattern that will reveal those real preference" (Sen, 1973, p. 252). Hausman offers a number of additional criticisms (1992). The empirical advantages of revealed preference theory, despite their motivating role for the theory, are meager. One proposal for determining whether an agent is indifferent between options is to see if the agent chooses each option roughly half the time. Determining whether an agent's preferences satisfy certain axioms of revealed preference requires observing a long sequence of repeated choices to decide whether the frequency of one option being chosen over another is close enough to attribute indifference or violation of the axioms. This is not even to raise the possibility of preference change during the long run of observations, or a taste for variety in the short run. "It is easier," Hausman claims, "to ask people what they prefer" (1992, p. 20). Revealed preference theory also trivializes certain seemingly substantial questions about rationality and choice. Should an agent's preferences be complete? They are typically defined to be in revealed preference theory. Do people sometimes choose against their preferences for, e.g., moral reasons? Not if choice reveals preference. Distinctions between choice and reflex movements cannot be marked, despite the subject matter being one of intentional, rational choice. Explaining or justifying choice in terms of preference becomes a trivial affair.

A relevant upshot of many of these considerations, it seems to me, is that care must be taken not to conflate issues having to do with aggregating preferences with certain issues having to do with choice, voting, and other behaviors. Against a revealed preference analysis, we can say that choice does not reveal preference and that higher-order criteria are not resolutions in value conflict at the first tier even if they determine a unique choice. Similarly, voting is a distinct issue from the problem of articulating compelling conceptions of consensus among some set of preference orderings. If we take choice to reveal preference, allowing tie-breaking or voting at the second (third, etc.) tier may entail violations of the Arrow axioms for social choice functions or social revealed preference.

But if choice does not reveal preference, those axioms may hold for *preference aggregation* even if they do not for social preference "revealed" by voting or some other choice procedure. My efforts are certainly more modest than Arrow's. I have attempted here to sketch an account of aggregating preferences, extending the Arrovian framework to cover aggregation with indeterminate preferences. The problem I set myself is to articulate a notion of neutrality or consensus for a given set of preference orderings. Whether the particular approach I take—suspend judgment among the candidate orderings—is found compelling or not in light of the offered motivations and possibility results, we should be careful not to prejudge questions about how preference relates to choice, or to assume constraints on preference aggregation are also reasonable constraints on voting.

Of course, people can and do vote with no eye at all to securing a consensus in preferences, as no part of a lexicographic choice procedure initiating from a consensus. The relevance of such voting and associated issues of strategic manipulation of the aggregation procedure to a normative account of consensus in preference is very far from clear. The task of articulating a neutral position for a given profile of preference orderings is distinct from the task of discerning whether preferences have been misrepresented or accurately measured or ensuring that individuals do not have incentives to misrepresent their preferences (in voting or similar situations) in light of knowledge of the aggregation rule. That an aggregation rule is open to manipulation would seem to have little to do with whether the rule provides a compelling account of consensus in preference. Voting may very well fail to reveal an agent's preferences between candidates when, say, voting for a less preferred option will lead to a more preferred outcome. (Estlund, for example, also distinguishes issues of voting from those of preference aggregation (2005, pp. 218-222).) As I suggested earlier, in bargaining, consensus in value judgments, preference, or the like may be of little concern. Each member of a group faces her own decision problem instead of facing a single group decision problem. That bargaining and voting are important features of social and political life is beyond dispute. What takes argument is the assimilation of consensus in preference to voting or bargaining.

9.2.3. *Dominance and Money Pumps.* Let's turn now to the second concern about this account of choice mentioned above. I claim that while the account does not reduce to binary comparisons between options, there is no difficulty in understanding the choice procedure for a set of orderings. But there is one argument that the account allows for *irrational* choices. Specifically, when we allow for second-tier (and higher) criteria of choice, revealed (basic) preference may be *cyclic*, even though all permissible orderings are individually acyclic.[14]

Example 3. *Consider the philosophy professor's decision again. In light of her professional goals, she orders the departments x over y over z. For considerations having to do with her private life, she ranks the departments y over z over x. Suppose now that she consults her old advisor who ranks the departments z over x over y for her. She decides to use this ranking as a tie-breaker. In that case, we have the following choices. $C(\{x,y\}) = \{x\}$, $C(\{y,z\}) = \{y\}$, and $C(\{x,z\}) = \{z\}$. The professor has revealed (strict, basic) preferences for x over y, for y over z, and for z over x.*

Davidson, McKinsey, and Suppes offered the "money pump argument" against cyclic preferences, and such arguments have been quite influential since (Davidson et al., 1955). Since the above preferences are strict, the professor should be willing to pay some (perhaps very small) fee to exchange option x for z. The fee could be a reduction in salary at z, for example. Similarly, she should be willing to pay at least a small fee to exchange z for y. But given the preference for x over y, she should be willing to pay a fee to return to her initial option, x, having lost money along the way.

One response to the money pump argument is that it fails to show any irrationality because it relies on *diachronic* choices. For example, Schick suggests that an agent might very well catch on to the scheme and refuse to be pumped (1986). Another way of putting the point is that an agent might prefer x to y, but prefer (a) y after exchanging x for z and z for y to (b) x after exchanging x for z and z for y. Fair enough. But Gustafsson thinks the diachronic element of the argument

[14] R is cyclic if for some $n \in \mathbb{Z}^+$ there are $x_1, ..., x_n \in X$ such that $x_1 P x_2, x_2 P x_3, ..., x_{n-1} P x_n, x_n P x_1$.

is beside the point, responding with a *synchronic* money pump argument (2013). By itself, losing money is not irrational. Though atypical, an agent could even *want* to lose money after all. What is irrational is choosing against one's preferences. Davidson, et al. diagnose the irrationality present in a money pump as a case of violating maximization: "a rational choice is one which selects an alternative to which none is preferred" (1955, p. 145). Accordingly, Gustafsson proposes doing away with any diachronic element of the argument and invites us to consider the single choice from the menu $\{x, y, z\}$. Gustafsson calls maximization "the non-dominated choice principle." He claims that cyclic preferences require an agent to select an option to which another is preferred. So any choice from the menu violates the non-dominated choice principle and is thereby irrational.

In the case of indeterminacy, however, admissibility does not coincide with revealed basic preference. In the three option menu, both x and y are admissible according to the professor's preferences. Applying the tie-breaker, x alone is left. Choice in the three option menu is not determined by the revealed preference relation for two option menus, otherwise the choice set would be empty. As Levi puts the point, "[r]evealed (revealed basic) preference is an epiphenomenon and not a fundamental determinant of the admissibility of options" (Levi, 1986a, p. 106). Moreover, choosing a dominated option is not licensed, at least not if dominance is to be defined in terms of categorical preference. We could say that an option $y \in Y$ is *dominated* if for every $R \in \mathcal{O}$ there is an $x \in Y$ such that xPy. In that case, y is not optimal according to any permissible ordering. But then y is not V-admissible in Y. Provided V-admissibility determines admissibility at every tier, no dominated option will ever be admissible. At the second tier, for example, a set of orderings will determine the V_2-admissible options among those V_1-admissible at the first tier. Finally, if we allow preference change, as both Dewey and deliberative democrats would insist we should, care must be taken with diachronic versions of money pump arguments. The risk is conflating preference change with cyclic preferences (Levi, 2002). An agent who prefers x to y now, but y to x in the future would be willing to pay a fee to exchange y for x today, and another fee to exchange x for y later.

Losing money in this fashion is not irrational. In this sense, money pump arguments prove too much, being insensitive to an important distinction.

9.3. Second Best vs. Second Worst and Informational Poverty. Consider the setup from Example 2 again. The menu consists of three options, $Y = \{x, y, z\}$, and $n = 2$ with xP_1yP_1z and that zP_2yP_2x so that R_1 and R_2 are in ordinal conflict. But suppose further that agent 1 feels that y is very close to x, but far from z, and agent 2 feels y is much closer to z than to x. The two agents agree, then, that y is "closer" to the best option. In such a case, y may seem a tempting social choice. If, on the other hand, both agents considered y closer to the worst option, it may seem tempting to exclude y from the social choice (Fishburn, 1973, p. 10). The relevant distinction here is between an option's being second best, and its being second worst (Levi, 1986a; Seidenfeld et al., 2010). The ability to draw the distinction requires *more than simple ordinal information* about an agent's preferences. It requires some *intrapersonal* preference intensity comparisons. Similarly, considerations of how much *stronger* i's preference for x over y is than j's preference for y over x requires more than ordinal information.

The point is well taken. It seems to me that such comparisons are routinely and sensibly made by decision makers. So, it is important for a theory of rational or social choice to respect the second best/second worst distinction. And that interpersonal comparisons are meaningful, despite contentions to the contrary from Arrow and others, finds fairly wide support. It may very well be the case that any reasonable form of deliberation relevant to the strengthening of consensus with which deliberative democrats are concerned requires quite strong interpersonal comparability assumptions. Like Fishburn, I consider the focus on preference orderings a self-imposed limitation of the present study and welcome further elaboration of the framework proposed in this chapter by "widening the informational basis" for collective choice. My dialectical strategy in accepting the restriction to ordinal information is to show the potential of indeterminacy even with informational poverty.

9.4. Generality. I have indicated that I think determinacy is overly restrictive for accounts of rational valuing even for individual agents. Yet, the framework presented here takes profiles of preference orderings as input. $\mathcal{F}$ could be interpreted as aggregating conflicting preferences for a single agent. But given the view that rationality does not require determinacy of preference, a full account of social choice with indeterminacy requires aggregating preferences which are themselves possibly indeterminate. In that case, additional complexities may emerge. Such an account requires further consideration, but is relegated to future work.

10. Conclusion

Given impossibility results like Arrow's and Sen's and the associated conceptual tensions involving deliberative democracy and non-welfarism, it should be clear that the classic, Arrovian social choice framework and axioms make heavy-handed assumptions about rational aggregation. The present essay explores a framework with a gentler touch. The central idea is that *sets* of preference orderings provide a compelling alternative conceptualization of consensus as shared agreement in preference. A set of preferences offers a neutral position from which to engage in group deliberation. Moreover, this conceptualization can be put to work. The axioms involved in Arrow's and Sen's results can be simultaneously satisfied. As a consequence, certain conceptual tensions in democratic theory and welfare economics dissolve. A suitably generalized social choice framework does not reveal "grievous foundational flaws" in democratic thought and practice. And, *pace* Riker, meaningful democratic choice is possible. The "tension between logicality and fairness" is not as "unresolvable" as it has been made out to be. Similarly, the possibility of satisfying prominent welfarist and procedural constraints demonstrated by Proposition 26 promises further resolution of conflicts between welfarist and non-welfarist concerns in welfare economics. I think that deliberative democrats should find something quite attractive in this view of social choice theory. They also have something to contribute to it. If deliberative democrats are right, the sorts of weak consensus positions identified above are not the final word on social preference and conflicts in value.

Bibliography

Aczél, J. (1948). On mean values. *Bulletin of the American Mathematical Society 54*(4), 392–400.

Aczél, J. and H. Oser (2006). *Lectures on Functional Equations and Their Applications*. Courier Corporation.

Aizerman, M. and A. Malishevski (1981). General theory of best variants choice: Some aspects. *Automatic Control, IEEE Transactions on 26*(5), 1030–1040.

Alchourrón, C. E., P. Gärdenfors, and D. Makinson (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of symbolic logic*, 510–530.

Arló-Costa, H. (2007). The logic of conditionals. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2014 ed.).

Arló-Costa, H. and J. Helzner (2010). Ambiguity aversion: The explanatory power of indeterminate probabilities. *Synthese 172*(1), 37–55.

Arló-Costa, H. and I. Levi (2006). Contraction: On the decision-theoretical origins of minimal change and entrenchment. *Synthese 152*(1), 129–154.

Arló-Costa, H. and H. Liu (2010). Value-based contraction: A representation result. Unpublished Manuscript. Carnegie Mellon University.

Arrow, K. J. (2012, originally published in 1951). *Social Choice and Individual Values*. Martino Publishing.

Bacharach, M. (1972). Scientific disagreement. *Unpublished Manuscript*.

Baratgin, J. and G. Politzer (2010). Updating: A psychologically basic situation of probability revision. *Thinking & Reasoning 16*(4), 253–287.

Belot, G. (2013). Bayesian orgulity. *Philosophy of Science 80*(4), 483–503.

Billingsley, P. (2008). *Probability and Measure*. John Wiley & Sons.

Black, D. (1958). *The Theory of Committees and Elections*. Springer.

Blackwell, D. and L. E. Dubins (1962). Merging of opinions with increasing information. *The Annals of Mathematical Statistics*, 882–886.

Blackwell, D. and L. E. Dubins (1975). On existence and non-existence of proper, regular, conditional distributions. *The Annals of Probability 3*(5), 741–752.

Bohman, J. (1997). *Deliberative Democracy: Essays on Reason and Politics*. MIT Press.

Booth, R. and J. Chandler (2016). The irreducibility of iterated to single revision. *Journal of Philosophical Logic*, Forthcoming.

Bordley, R. F. (1982). A multiplicative formula for aggregating probability assessments. *Management Science 28*(10), 1137–1148.

Bossert, W. and M. Fleurbaey (2015). An interview with kotaro suzumura. *Social Choice and Welfare 44*(1), 179.

Boutilier, C. (1993). Revision sequences and nested conditionals. In *IJCAI*, Volume 93, pp. 519–531.

Bradley, S. (2014). Imprecise probabilities. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2014 ed.).

Buchanan, J. M. (1954). Social choice, democracy, and free markets. *The Journal of Political Economy*, 114–123.

Buchanan, J. M. and G. Tullock (1962). *The Calculus of Consent*, Volume 3. University of Michigan Press Ann Arbor.

Chang, R. (2009). Reflections on the reasonable and the rational in conflict resolution. *Aristotelian Society Supplementary Volume 83*(1), 133–160.

Christensen, D. (2009). Disagreement as evidence: The epistemology of controversy. *Philosophy Compass 4*(5), 756–767.

Cisewski, J., J. Kadane, M. Schervish, T. Seidenfeld, and R. Stern (MS). Standards for *Modest* bayesian credenes. *Philosophy of Science*, Forthcoming.

Clemen, R. T. and R. L. Winkler (1999). Combining probability distributions from experts in risk analysis. *Risk Analysis 19*(2), 187–203.

Cohen, J. (1997). Deliberation and democratic legitimacy. In *Deliberative Democracy: Essays on Reason and Politics*, pp. 67. MIT Press.

Cozman, F. G. (1998). Irrelevance and independence relations in quasi-bayesian networks. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, pp. 89–96. Morgan Kaufmann Publishers Inc.

Cozman, F. G. (2000). Credal networks. *Artificial Intelligence 120*(2), 199–233.

Dahl, R. A. (1989). *Democracy and Its Critics.* Yale University Press.

Darwiche, A. and J. Pearl (1997). On the logic of iterated belief revision. *Artificial Intelligence 89*(1), 1–29.

Davidson, D., J. C. C. McKinsey, and P. Suppes (1955). Outlines of a formal theory of value, i. *Philosophy of Science 22*(2), 140–160.

de Campos, L. M. and S. Moral (1995). Independence concepts for convex sets of probabilities. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 108–115. Morgan Kaufmann Publishers Inc.

de Finetti, B. (1931). Sul concetto di media. *Giornale dell'Istituto Italiano degli Attuari 2*, 369–396.

de Finetti, B. (1964). Foresight: Its logical laws, its subjective sources. In H. E. Kyburg and H. E. Smoklery (Eds.), *Studies in Subjective Probability.* Wiley.

de Finetti, B. (1990). *Theory of Probability*, Volume 1. John Wiley and Sons.

Dewey, J. (1941). Propositions, warranted assertibility, and truth. *The Journal of Philosophy 38*(7), 169–186.

Dewey, J. and J. H. Tufts (1932). Ethics (revised edition). *New York: H. Holt and Company.*

Diaconis, P. and S. L. Zabell (1982). Updating subjective probability. *Journal of the American Statistical Association 77*(380), 822–830.

Dietrich, F. (2010). Bayesian group belief. *Social Choice and Welfare 35*(4), 595–626.

Dietrich, F. and C. List (2014). Probabilistic opinion pooling. In A. Hájek and C. Hitchcock (Eds.), *Oxford Handbook of Probability and Philosophy*. Oxford University Press.

Dietrich, F. and C. List (2016). Mentalism versus behaviourism in economics: a philosophy-of-science perspective. *Economics and Philosophy 32*(02), 249–281.

Dryzek, J. S. and C. List (2003). Social choice theory and deliberative democracy: a reconciliation. *British Journal of Political Science 33*(01), 1–28.

Duggan, J. (1999). A general extension theorem for binary relations. *Journal of Economic Theory 86*(1), 1–16.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Elga, A. (2007). Reflection and disagreement. *Noûs 41*(3), 478–502.

Elkin, L. (2015). An epistemically modest response to disagreement, agm-ified. *The Reasoner 9*(9), 76–77.

Elkin, L. and G. Wheeler (2016). Resolving peer disagreements through imprecise probabilities. *Noûs*, DOI: 10.1111/nous.12143.

Ellsberg, D. (1963). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics 77*(2), 327–336.

Elster, J. (1997). The market and the forum: Three varieties of political theory. In *Deliberative Democracy: Essays on Reason and Politics*, pp. 3–34. MIT Press Cambridge, MA.

Elster, J. (2007). *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. Cambridge University Press.

Estlund, D. (2005). Democratic theory. In *Oxford Handbook of Contemporary Philosophy*, pp. 208–30. Oxford University Press.

Feldman, R. (2011). Reasonable religious disagreements. In A. Goldman and D. Whitcomb (Eds.), *Social Epistemology: Essential Readings*, Volume 137. Oxford University Press.

Field, H. (1978). A note on jeffrey conditionalization. *Philosophy of Science*, 361–367.

Fishburn, P. C. (1973). *The Theory of Social Choice*, Volume 264. Princeton University Press Princeton.

French, S. (1985, September). Group consensus probability distributions: A critical survey. In D. L. J.M. Bernardo, M.H. DeGroot and A. Smith (Eds.), *Bayesian Statistics: Proceedings of the Second Valencia International Meeting*, Volume 2, pp. 183–201. North-Holland.

Fuhrmann, A. and S. O. Hansson (1994). A survey of multiple contractions. *Journal of Logic, Language and Information 3*(1), 39–75.

Gaifman, H. and M. Snir (1982). Probabilities over rich languages, testing and randomness. *The Journal of Symbolic Logic 47*(03), 495–548.

Gaifman, H. and A. Vasudevan (2012). Deceptive updating and minimal information methods. *Synthese 187*(1), 147–178.

Gärdenfors, P. (1982). Imaging and conditionalization. *The Journal of Philosophy*, 747–760.

Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA.

Gärdenfors, P. and H. Rott (1995). *Handbook of Logic in Artificial Intelligence and Logic Programming: Epistemic and temporal reasoning*, Volume 4, Chapter Belief Revision. Oxford University Press, Oxford.

Gärdenfors, P. and N.-E. Sahlin (1982). Unreliable probabilities, risk taking, and decision making. *Synthese 53*(3), 361–386.

Genest, C. (1984). A characterization theorem for externally bayesian groups. *The Annals of Statistics*, 1100–1105.

Genest, C., K. J. McConway, and M. J. Schervish (1986). Characterization of externally bayesian pooling operators. *The Annals of Statistics*, 487–501.

Genest, C. and C. G. Wagner (1987). Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae 32*(1), 74–86.

Genest, C. and J. V. Zidek (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 114–135.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Girón, F. J. and S. Ríos (1980). Quasi-bayesian behaviour: A more realistic approach to decision making? *Trabajos de Estadística y de Investigación Operativa 31*(1), 17–38.

Good, I. J. (1983). *Good Thinking: The Foundations of Probability and Its Applications.* U of Minnesota Press.

Guignon, C. and D. R. Hiley (2003). Richard rorty and contemporary philosophy. In *Richard Rorty.* Cambridge University Press.

Gustafsson, J. E. (2013, August). The irrelevance of the diachronic money-pump argument for acyclicity. *The Journal of Philosophy 110*(8), 460–464.

Hájek, A. and N. Hall (1994). The hypothesis of the conditional construal of conditional probability. In E. Eells and B. Skyrms (Eds.), *Probability and Conditionals: Belief Revision and Rational Decision*, pp. 75–112. Cambridge University Press.

Halmos, P. R. (1963). *Lectures on Boolean Algebras.* Van Nostrand, Princeton.

Hansson, B. (1968). Choice structures and preference relations. *Synthese 18*(4), 443–458.

Hansson, S. O. (1993). Money-pumps, self-torturers and the demons of real life. *Australasian Journal of Philosophy 71*(4), 476–485.

Hansson, S. O. (1999). *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Volume 11. Springer Science & Business Media.

Hansson, S. O. (2003). Ten philosophical problems in belief revision. *Journal of Logic and Computation 13*(1), 37–49.

Hansson, S. O. and E. J. Olsson (1995). Levi contractions and agm contractions: A comparison. *Notre Dame Journal of Formal Logic 36*(1), 103–119.

Hardin, R. (1993). Public choice versus democracy. *The Idea of Democracy*, 157–72.

Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy 61*(5), 434–435.

Hartmann, S. (2014). A new solution to the problem of old evidence. In *Philosophy of Science Association 24th Biennial Meeting*, Chicago, IL.

Hausman, D. M. (1992). *The Inexact and Separate Science of Economics.* Cambridge University Press.

Helzner, J. (2013). Rationalizing two-tiered choice functions through conditional choice. *Synthese 190*(6), 929–951.

Herron, T., T. Seidenfeld, and L. Wasserman (1997). Divisive conditioning: Further results on dilation. *Philosophy of Science*, 411–444.

Herzberg, F. (2014). Aggregating infinitely many probability measures. *Theory and Decision*, 1–19.

Huttegger, S. M. (2013). In defense of reflection. *Philosophy of Science 80*(3), 413–433.

Huttegger, S. M. (2015a). Bayesian convergence to the truth and the metaphysics of possible worlds. *Philosophy of Science 82*(4), 587–601.

Huttegger, S. M. (2015b). Merging of opinions and probability kinematics. *The Review of Symbolic Logic 8*(04), 611–648.

Jeffrey, R. (2004). *Subjective Probability: The Real Thing.* Cambridge University Press.

Joyce, J. M. (1999). *The Foundations of Causal Decision Theory.* Cambridge University Press.

Kalai, E. and E. Lehrer (1994). Weak and strong merging of opinions. *Journal of Mathematical Economics 23*(1), 73–86.

Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge University Press.

Keeney, R. L. and H. Raiffa (1993). *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press.

Kelly, K. T. (1996). *The Logic of Reliable Inquiry*. Oxford University Press.

Kolmogorov, A. N. (1930). Sur la notion de la moyenne. *Atti della R. Academia Nazionale dei Lincei 12*(9), 388–391.

Kullback, S. and R. A. Leibler (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 79–86.

Kyburg, H. E. (1987). Bayesian and non-bayesian evidential updating. *Artificial Intelligence 31*(3), 271–293.

Kyburg, H. E. (1998). Interval-valued probabilities. *Imprecise Probabilities Project*.

Kyburg, H. E. and M. Pittarelli (1992). Some problems for convex bayesians. In *Proceedings of the Eighth international conference on Uncertainty in artificial intelligence*, pp. 149–154. Morgan Kaufmann Publishers Inc.

Kyburg, H. E. and M. Pittarelli (1996). Set-based bayesianism. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on 26*(3), 324–339.

Larmore, C. (2003). Public reason. In S. R. Freeman (Ed.), *The Cambridge Companion to Rawls*, pp. 368–393. Cambridge University Press.

Lehrer, K. and C. Wagner (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*, Volume 21. Springer.

Lehrer, K. and C. Wagner (1983). Probability amalgamation and the independence issue: A reply to laddaga. *Synthese 55*(3), 339–346.

Leitgeb, H. (2016). Imaging all the people. *Episteme* (DOI: 10.1017/epi.2016.14).

Levi, I. (1967a). Gambling with truth: An essay on induction and the aims of science.

Levi, I. (1967b). Probability kinematics. *British Journal for the Philosophy of Science*, 197–209.

Levi, I. (1970). Probability and evidence. In M. Swain (Ed.), *Induction, Acceptance, and Rational Belief*, pp. 134–156. New York: Humanities Press.

Levi, I. (1974). On indeterminate probabilities. *The Journal of Philosophy 71*(13), 391–418.

Levi, I. (1978). Irrelevance. In C. Hooker, J. Leach, and E. McClennen (Eds.), *Foundations and Applications of Decision Theory*, Volume 1, pp. 263–273. Boston: Springer.

Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press, Cambridge, MA.

Levi, I. (1982). Conflict and social agency. *The Journal of Philosophy*, 231–247.

Levi, I. (1985). Consensus as shared agreement and outcome of inquiry. *Synthese 62*(1), pp. 3–11.

Levi, I. (1986a). *Hard choices: Decision making under unresolved conflict*. Cambridge University Press.

Levi, I. (1986b). The paradoxes of allais and ellsberg. *Economics and Philosophy 2*(1), 23–53.

Levi, I. (1990). Pareto unanimity and consensus. *The Journal of Philosophy 87*(9), 481–492.

Levi, I. (1991). *The Fixation of Belief and Its Undoing: Changing Beliefs through Inquiry*. Cambridge University Press.

Levi, I. (1996). *For the Sake of the Argument: Ramsey Test Conditionals, Inductive Inference and Nonmonotonic Reasoning*. Cambridge University Press.

Levi, I. (1998). Pragmatism and change of view. *Canadian Journal of Philosophy 28*(sup1), 177–201.

Levi, I. (2002). Money pumps and diachronic books. *Philosophy of Science 69*(S3), S235–S247.

Levi, I. (2004). *Mild Contraction: Evaluating Loss of Information Due to Loss of Belief*. Oxford University Press.

Levi, I. (2006). Escape from boredom: Edification according to rorty. In *Decisions and Revisions: Philosophical Essays on Knowledge and Value*, pp. 136–146. Cambridge University Press.

Levi, I. (2008). Convexity and separability in representing consensus. In K. Basu and R. Kanbur (Eds.), *Arguments for a Better World: Essays in Honor of Amartya Sen*, Volume 1: Ethics, Welfare, and Measurement, Chapter 11, pp. 193–212. Oxford University Press.

Levi, I. (2009). Why indeterminate probability is rational. *Journal of Applied Logic 7*(4), 364–376.

Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *The Philosophical Review 85*, 297–315.

List, C. (2013). Social choice theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2013 ed.).

List, C. and P. Pettit (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press Oxford.

Mackie, G. (2003). *Democracy Defended*. Cambridge University Press.

Madansky, A. (1964). Externally bayesian groups. Santa Monica, CA: RAND Corporation.

McConway, K. J. (1981). Marginalization and linear opinion pools. *Journal of the American Statistical Association 76*(374), 410–414.

Miller, D. (1992). Deliberative democracy and social choice. *Political studies 40*(1 suppl), 54–67.

Misak, C. (2004). *Truth and the End of Inquiry: A Peircian Account of Truth*. Oxford: Clarendon Press.

Mongin, P. (1995). Consistent bayesian aggregation. *Journal of Economic Theory 66*(2), 313–351.

Mongin, P. (1997). Spurious unanimity and the pareto principle. *Economics and Philosophy*, 1–22.

Moral, S. and J. Del Sagrado (1998). Aggregation of imprecise probabilities. In B. Bouchon-Meunier (Ed.), *Aggregation and Fusion of Imperfect Information*, pp. 162–188. Springer.

Moulin, H. (1985). Choice functions over a finite set: a summary. *Social Choice and Welfare 2*(2), 147–160.

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

Nayak, A. C., M. Pagnucco, and P. Peppas (2003). Dynamic belief revision operators. *Artificial Intelligence 146*(2), 193–228.

Nielsen, M. (MS). Deterministic convergence and strong regularity. Unpublished Manuscript.

Ouchi, F. (2004). A literature review on the use of expert opinion in probabilistic risk analysis.

Pedersen, A. P. (2009a). Pseudo-rationalizability over infinite choice spaces. Technical report, Citeseer.

Pedersen, A. P. (2009b). Rational choice and formal epistemology. Master's thesis, Carnegie Mellon University.

Pedersen, A. P. and G. Wheeler (2014). Demystifying dilation. *Erkenntnis 79*(6), 1305–1342.

Pedersen, A. P. and G. Wheeler (2015). Dilation, disintegrations, and delayed decisions. In *Proceedings of the 9th*.

Peirce, C. S. (1992a). *The Essential Peirce: Selected Philosophical Writings*, Volume 1. Indiana University Press.

Peirce, C. S. (1992b). How to make our ideas clear. In N. Houser and C. Kloesel (Eds.), *The Essential Peirce, Volume 1: Selected Philosophical Writings (1867–1893)*, pp. 124–141. Indiana University Press.

Perote-Peña, J. and A. Piggins (2015). A model of deliberative and aggregative democracy. *Economics and Philosophy 31*(1), 93.

Quine, W. V. O. (2013). *Word and Object.* MIT Press.

Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Random House.

Ramsey, F. P. (1990a). General propositions and causality. In H. Mellor (Ed.), *Philosophical Papers.* Cambridge University Press.

Ramsey, F. P. (1990b). Truth and probability. In D. H. Mellor (Ed.), *Philosophical Papers*, pp. 52–109. Cambridge University Press.

Redacted. Redacted.

Riker, W. (1982). *Liberalism Against Populism.* Waveland Press, Inc.

Rockafellar, R. T. (1970). *Convex Analysis.* Number 28. Princeton University Press.

Rorty, R. (1980). *Philosophy and the Mirror of Nature.* Cambridge University Press.

Rott, H. (2000). Two dogmas of belief revision. *The Journal of Philosophy*, 503–522.

Rott, H. and M. Pagnucco (1999). Severe withdrawal (and recovery). *Journal of Philosophical Logic 28*(5), 501–547.

Rovane, C. (1997). *The Bounds of Agency: An Essay in Revisionary Metaphysics.* Princeton University Press.

Russell, J. S., J. Hawthorne, and L. Buchak (2015). Groupthink. *Philosophical Studies 172*(5), 1287–1309.

Savage, L. (1972, originally published in 1954). *The Foundations of Statistics.* New York: John Wiley and Sons.

Schervish, M. and T. Seidenfeld (1990). An approach to consensus and certainty with increasing evidence. *Journal of Statistical Planning and Inference 25*(3), 401–414.

Schick, F. (1986). Dutch bookies and money pumps. *The Journal of Philosophy 83*(2), 112–119.

Seidenfeld, T. (1986). Entropy and uncertainty. *Philosophy of Science*, 467–491.

Seidenfeld, T. (1993). Outline of a theory of partially ordered preferences. *Philosophical Topics 21*(1), 173–189.

Seidenfeld, T. (2001). Remarks on the theory of conditional probability: Some issues of finite versus countable additivity. In V. F. Hendricks (Ed.), *Probability Theory: Philosophy, Recent History and Relations to Science.* Synthese Library, Kluwer.

Seidenfeld, T., J. B. Kadane, and M. J. Schervish (1989). On the shared preferences of two bayesian decision makers. *The Journal of Philosophy 86*(5), 225–244.

Seidenfeld, T., M. J. Schervish, and J. B. Kadane (2010). Coherent choice functions under uncertainty. *Synthese 172*(1), 157–176.

Seidenfeld, T. and L. Wasserman (1993). Dilation for sets of probabilities. *The Annals of Statistics 21*(3), 1139–1154.

Sen, A. (1970a). *Collective Choice and Social Welfare*. Holden Day.

Sen, A. (1970b). The impossibility of a paretian liberal. *Journal of Political Economy 78*(1), 152–157.

Sen, A. (1971). Choice functions and revealed preference. *The Review of Economic Studies 38*(3), 307–317.

Sen, A. (1973). Behaviour and the concept of preference. *Economica 40*(159), 241–259.

Sen, A. (1977). Social choice theory: A re-examination. *Econometrica: Journal of the Econometric Society*, 53–89.

Sen, A. (1997). Maximization and the act of choice. *Econometrica: Journal of the Econometric Society*, 745–779.

Sen, A. (2004a). The possibility of social choice. In *Rationality and Freedom.* Harvard University Press.

Sen, A. (2004b). *Rationality and Freedom.* Harvard University Press.

Skyrms, B. (1986). *Choice and Chance: An Introduction to Inductive Logic* (3rd Edition ed.). Belmont: Wadsworth Publishing Company.

Skyrms, B. (1996). The structure of radical probabilism. *Erkenntnis 45*(2-3), 285–297.

Smith, C. A. B. (1961). Consistency in statistical inference and decision. *Journal of the Royal Statistical Society. Series B (Methodological) 23*(1), pp. 1–37.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In W. L. Harper and B. Skyrms (Eds.), *Causation in Decision, Belief Change, and Statistics*, Volume 42 of *The University of Western Ontario Series in Philosophy of Science*, pp. 105–134. Kluwer

Academic Publishers.

Spohn, W. (2012). *The Laws of Belief: Ranking Theory and Its Philosophical Applications.* Oxford University Press, USA.

Stewart, R. T. and I. Ojea Quintana (2017a). Learning and pooling, pooling and learning. *Erkenntnis*, Forthcoming.

Stewart, R. T. and I. Ojea Quintana (2017b). Probabilistic opinion pooling with imprecise probabilities. *Journal of Philosophical Logic*, doi:10.1007/s10992–016–9415–9.

Stone, M. (1961). The opinion pool. *The Annals of Mathematical Statistics 32*(4), 1339–1342.

Suzumura, K. (1983). *Rational choice, Collective decisions, and Social Welfare.* Cambridge [etc.]: Cambridge University Press.

Szpilrajn, E. (1930). Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae 16*(1), 386–389.

van Fraassen, B. C. (1989). *Laws and Symmetry.* Clarendon Press Oxford.

Wagner, C. (2002). Probability kinematics and commutativity. *Philosophy of Science 69*(2), 266–278.

Wagner, C. (2009). Jeffrey conditioning and external bayesianity. *Logic Journal of IGPL 18*(2), 336–345.

Walley, P. (1982). The elicitation and aggregation of beliefs. Technical Report 23, Department of Statistics, University of Warwick, Coventry CV4 7AL, England.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities.* Chapman and Hall London.

Wasserman, L. (1993). Review: Statistical reasoning with imprecise probabilities by peter walley. *Journal of the American Statistical Association 88*(422), pp. 700–702.

Wasserman, L. and T. Seidenfeld (1994). The dilation phenomenon in robust bayesian inference. *Journal of Statistical Planning and Inference 40*(2), 345–356.

Williams, P. M. (1980). Bayesian conditionalisation and the principle of minimum information. *British Journal for the Philosophy of Science*, 131–144.

Appendices

APPENDIX A

Appendix to Chapter 2: Proofs

Proof of Proposition 1

Proof. We carry out McConway's proof with minimal adjustments made for our framework (pp. 411-412 1981, Theorem 3.1).

WSFP $\Rightarrow$ MP. Assume that $\mathcal{F}$ has the WSFP, i.e., there is a function $\mathcal{G} : \mathscr{A} \times [0,1]^n \to \mathscr{P}([0,1])$ such that $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{G}(A, \boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A))$. By WSFP, we have $\mathcal{F}([\boldsymbol{p}_1 \restriction_{\mathscr{A}'}], ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}])(A)$ $= \mathcal{G}(A, [\boldsymbol{p}_1 \restriction_{\mathscr{A}'}](A), ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}](A))$. Since $\mathcal{G}$ is a function and $\boldsymbol{p}_i(A) = [\boldsymbol{p}_i \restriction_{\mathscr{A}'}](A)$ for any $A \in \mathscr{A}'$ (all such $A \in \mathscr{A}'$ are also in $\mathscr{A}$), it follows that $\mathcal{G}(A, [\boldsymbol{p}_1 \restriction_{\mathscr{A}'}](A), ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}](A)) = \mathcal{G}(A, \boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A))$ $= \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$. Hence, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}([\boldsymbol{p}_1 \restriction_{\mathscr{A}'}], ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}])(A)$.

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

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

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

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}([\boldsymbol{p}_1 \restriction_{\mathscr{A}'}], ..., [\boldsymbol{p}_n \restriction_{\mathscr{A}'}])(A).$$

$\mathscr{A}'$ is uniquely defined by A and any probability over $\mathscr{A}'$ is uniquely determined by the probability of A under that distribution. So the righthand side of the equation above is determined by A and $\boldsymbol{p}_i \restriction_{\mathscr{A}'}(A) = [\boldsymbol{p}_i \restriction_{\mathscr{A}'}](A) = \boldsymbol{p}_i(A)$. □

Proof of Lemma 1

PROOF. Let $Y = \{\boldsymbol{p} : \boldsymbol{p} = \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i$ such that $\alpha_i \geqslant 0$ for $i = 1, ..., n$ and $\sum_{i=1}^{n} \alpha_i = 1\}$. We want to show the following:

$$\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv\{\boldsymbol{p}_i : i = 1, ..., n\} = Y$$

The first equality we have by definition. In order to show the second equality, we have to show that Y is the smallest convex set containing $\{\boldsymbol{p}_i : i = 1, ..., n\}$. To show convexity, we show that for any two functions in Y, any convex combination of those functions is in Y. Suppose that $\boldsymbol{p}, \boldsymbol{p}' \in Y$. By assumption, $\boldsymbol{p} = \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i$ and $\boldsymbol{p}' = \sum_{i=1}^{n} \beta_i \boldsymbol{p}_i$. Consider $\boldsymbol{p}^\star = \gamma \boldsymbol{p} + (1-\gamma)\boldsymbol{p}' = \gamma(\sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i) + (1-\gamma)\sum_{i=1}^{n} \beta_i \boldsymbol{p}_i$.

$$\begin{aligned}
\boldsymbol{p}^\star &= \gamma \textstyle\sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i + (1-\gamma)\sum_{i=1}^{n} \beta_i \boldsymbol{p}_i \\
&= \textstyle\sum_{i=1}^{n} \gamma \alpha_i \boldsymbol{p}_i + \sum_{i=1}^{n} (1-\gamma)\beta_i \boldsymbol{p}_i \\
&= \textstyle\sum_{i=1}^{n} [\gamma \alpha_i \boldsymbol{p}_i + (1-\gamma)\beta_i \boldsymbol{p}_i] \\
&= \textstyle\sum_{i=1}^{n} [\gamma \alpha_i + (1-\gamma)\beta_i] \boldsymbol{p}_i \\
&= \textstyle\sum_{j=1}^{n} \delta_j \boldsymbol{p}_j
\end{aligned}$$

where $\delta_j = \gamma \alpha_j + (1-\gamma)\beta_j$. $\delta_j \geqslant 0$ for $j = 1, ..., n$ because every term is nonnegative. $\sum_{j=1}^{n} \delta_j = \sum_{i=1}^{n} [\gamma \alpha_i + (1-\gamma)\beta_i] = \sum_{i=1}^{n} \gamma \alpha_i + \sum_{i=1}^{n} (1-\gamma)\beta_i = \gamma \sum_{i=1}^{n} \alpha_i + (1-\gamma)\sum_{i=1}^{n} \beta_i = \gamma(1) + (1-\gamma)1 = 1$. Hence, $\boldsymbol{p}^\star \in Y$, so Y is convex. If Y were not the smallest such set, then there would be some convex $Z \subsetneq Y$ such that $\{\boldsymbol{p}_i : i = 1, ..., n\} \subseteq Z$. But for any $\boldsymbol{p} \in Y, \boldsymbol{p}$ is a convex combination of the elements in $\{\boldsymbol{p}_i : i = 1, ..., n\}$. Since Z is convex and contains the $\boldsymbol{p}_i$, it follows that $\boldsymbol{p} \in Z$, which is a contradiction. □

Proof of Proposition 2

PROOF. Since $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$, we let $\mathcal{G}$ of the SSFP be the convex hull operation applied to $\{\boldsymbol{p}_i(A) : i = 1, ..., n\}$. It is clear that $\mathcal{G}$ depends just on the individual probabilities for A. We need to show that

$$\{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\} = conv\{\boldsymbol{p}_i(A) : i = 1, ..., n\}.$$

Trivially, the lefthand side includes $\{\boldsymbol{p}_i(A) : i = 1, ..., n\}$. Suppose $\boldsymbol{p}(A), \boldsymbol{p}'(A) \in \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$. Since $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is convex, it follows immediately that any convex combination of $\boldsymbol{p}(A), \boldsymbol{p}'(A)$ is in $\{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$. Finally, suppose that there is some convex $Z \subsetneq \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$ which contains $\{\boldsymbol{p}_i(A) : i = 1, ..., n\}$. But for any $\boldsymbol{p}(A) \in \{\boldsymbol{p}(A) : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$, $\boldsymbol{p}(A)$ is a convex combination of the $\boldsymbol{p}_i(A)$ since every $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is such a convex combination of the $\boldsymbol{p}_i$ (Lemma 1). Hence, $\boldsymbol{p}(A) \in Z$, contrary to our supposition. So, the equality holds and the SWFP is satisfied.

But since SWFP clearly implies WSFP, WSFP is satisfied, too. By Proposition 1, it follows immediately that $\mathcal{F}$ has the MP.

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

$$\begin{aligned}\{0\} &= \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(\varnothing)\\ &= \mathcal{G}(\boldsymbol{p}_1(\varnothing), ..., \boldsymbol{p}_n(\varnothing))\\ &= \mathcal{G}(\boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A))\\ &= \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)\end{aligned}$$

So, ZPP follows from SWFP.

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

□

Proof of Lemma 2

We generalize a proof of a result due originally to Girón and Rios and Levi (Levi, 1978; Girón and Ríos, 1980) for updating on an *event* to updating on a common likelihood function.

PROOF. We want to show that $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is convex. That is, given any two members, $\boldsymbol{p}^\lambda, \boldsymbol{p}'^\lambda \in \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ and $\alpha \in [0,1]$, $\boldsymbol{p}^\star = \alpha \boldsymbol{p}^\lambda + (1-\alpha)\boldsymbol{p}'^\lambda$ is in $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. If there is a convex combination of $\boldsymbol{p}$ and $\boldsymbol{p}'$, $\boldsymbol{p}_\star$, such that $\boldsymbol{p}_\star^\lambda = \boldsymbol{p}^\star$, then the convexity of $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is established as a consequence of the convexity of $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Where $\boldsymbol{p}_\star^\lambda(\cdot) = \frac{\boldsymbol{p}_\star(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}_\star(\omega')\lambda(\omega')} = \frac{\beta \boldsymbol{p}(\cdot)\lambda(\cdot) + (1-\beta)\boldsymbol{p}'(\cdot)\lambda(\cdot)}{\beta \sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega') + (1-\beta)\sum_{\omega' \in \Omega} \boldsymbol{p}'(\omega')\lambda(\omega')}$, for any α we want to find some β such that

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

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

Proof of Proposition 3

PROOF. We must show that convex IP pooling functions are externally Bayesian, i.e., $\mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) = \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ (provided the relevant profiles are in the domain of $\mathcal{F}$).

$\mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) \subseteq \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Trivially, for each $i = 1, ..., n$, $\boldsymbol{p}_i^\lambda \in \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. By Lemma 2, $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is convex. It follows that $conv\{\boldsymbol{p}_i^\lambda : i = 1, ..., n\} = \mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) \subseteq \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$.

$\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \subseteq \mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda)$. By Lemma 1, any $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ can be expressed as the convex combination of the n extreme points generating $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, i.e., $\boldsymbol{p} = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i$ where $\alpha_i \geqslant 0$ for

$i = 1, ..., n$ and $\sum_{i=1}^{n} \alpha_i = 1$. By definition,

$$\mathcal{F}^{\lambda}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \{\boldsymbol{p}^{\lambda} : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \text{ and } \boldsymbol{p}^{\lambda}(\cdot) = \frac{\boldsymbol{p}(\cdot)\lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')}\}$$

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

$$\begin{aligned}
\boldsymbol{p}^{\lambda}(\omega) &= \frac{\boldsymbol{p}(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \boldsymbol{p}(\omega')\lambda(\omega')} && \text{[Definition]} \\
&= \frac{\sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i(\omega)\lambda(\omega)}{\sum_{\omega' \in \Omega} \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i(\omega')\lambda(\omega')} && \text{[Lemma 1]} \\
&= \frac{\sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i^{\lambda}(\omega) \cdot \sum_{\omega' \in \Omega} \boldsymbol{p}_i(\omega')\lambda(\omega')}{\sum_{\omega' \in \Omega} \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i(\omega')\lambda(\omega')} && [\boldsymbol{p}_i(\omega)\lambda(\omega) = \boldsymbol{p}_i(\omega)^{\lambda} \cdot \textstyle\sum_{\omega' \in \Omega} \boldsymbol{p}_i(\omega')\lambda(\omega')] \\
&= \sum_{j=1}^{n} \beta_j \boldsymbol{p}_j^{\lambda}(\omega) \in \mathcal{F}(\boldsymbol{p}_1^{\lambda}, ..., \boldsymbol{p}_n^{\lambda}) && \text{[Algebra]}
\end{aligned}$$

where $\beta_j = \dfrac{\alpha_j \cdot \sum_{\omega' \in \Omega} \boldsymbol{p}_j(\omega')\lambda(\omega')}{\sum_{\omega' \in \Omega} \sum_{i=1}^{n} \alpha_i \boldsymbol{p}_i(\omega')\lambda(\omega')}$ with $\beta_j \geqslant 0$ for all $j = 1, ..., n$ and $\sum_{j=1}^{n} \beta_j = 1$.

□

Proof of Proposition 4

Proof. We provide a very simple type of counterexample to individualwise Bayesianity, though counterexamples are plentiful. Consider the profile $(\boldsymbol{p}_1, \boldsymbol{p}_2)$ for $n = 2$ agents such that $\boldsymbol{p}_1 = \boldsymbol{p}_2$. Individualwise Bayesianity requires that $\mathcal{F}(\boldsymbol{p}_1, \boldsymbol{p}_2^{\lambda}) = \mathcal{F}^{\lambda}(\boldsymbol{p}_1, \boldsymbol{p}_2)$ (provided both $(\boldsymbol{p}_1, \boldsymbol{p}_2)$ and $(\boldsymbol{p}_1, \boldsymbol{p}_2^{\lambda})$ are in the domain of $\mathcal{F}$). By Proposition 3 (external Bayesianity), it follows that $\mathcal{F}^{\lambda}(\boldsymbol{p}_1, \boldsymbol{p}_2) = \mathcal{F}(\boldsymbol{p}_1^{\lambda}, \boldsymbol{p}_2^{\lambda})$. But since $\boldsymbol{p}_1 = \boldsymbol{p}_2$, it follows that $\boldsymbol{p}_1^{\lambda} = \boldsymbol{p}_2^{\lambda}$. By unanimity (Proposition 2), then, we have $\mathcal{F}(\boldsymbol{p}_1^{\lambda}, \boldsymbol{p}_2^{\lambda}) = \{\boldsymbol{p_i}^{\lambda}\}$, where $\boldsymbol{p}_i^{\lambda} = \boldsymbol{p}_1^{\lambda} = \boldsymbol{p}_2^{\lambda}$. However, in general $\boldsymbol{p}_i \neq \boldsymbol{p}_i^{\lambda}$ and so $\mathcal{F}(\boldsymbol{p}_1, \boldsymbol{p}_2^{\lambda})$ is *not* a singleton. It follows that, in general, $\mathcal{F}(\boldsymbol{p}_1, \boldsymbol{p}_2^{\lambda}) \neq \mathcal{F}^{\lambda}(\boldsymbol{p}_1, \boldsymbol{p}_2)$.

□

Proof of Proposition 5

Proof. Suppose that $\boldsymbol{p}_i(A|B) = \boldsymbol{p}_i(A)$ for $i = 1, ..., n$. We want to show that $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$ $= \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$. Consider $\boldsymbol{p}^\star(A) \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$ and $\boldsymbol{p}_\star(A) \in \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$. By Lemma 1, $\boldsymbol{p}^\star(A) = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i(A)$, for appropriate α_i. By Proposition 3 (external Bayesianity), $\mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}(\boldsymbol{p}_1^B, ..., \boldsymbol{p}_n^B)(A)$ (Proposition 3 holds for standard conditionalization since standard conditionalization is a special case of updating on a likelihood function, as noted in the body of the essay). So, we have $\boldsymbol{p}_\star(A) = \sum_{i=1}^n \beta_i \boldsymbol{p}_i^B(A)$, for appropriate β_i, again by Lemma 1. By hypothesis $\boldsymbol{p}_i^B(A) = \boldsymbol{p}_i(A)$ for $i = 1, ..., n$. Hence, $\boldsymbol{p}_\star(A) = \sum_{i=1}^n \beta_i \boldsymbol{p}_i^B(A) = \sum_{i=1}^n \beta_i \boldsymbol{p}_i(A)$. Letting $\alpha_i = \beta_i$, it follows that $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A) = \mathcal{F}^B(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)(A)$. □

Proof of Proposition 6

Proof. We show first that $\mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \subseteq conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ for all $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$. Since there are at least three disjoint events, $A_1, A_2, A_3 \in \mathscr{A}$, following (Lehrer and Wagner, 1981, Theorems 6.4, 6.7) and (McConway, 1981, Theorem 3.3), we can exploit techniques and results for functional equations. For any numbers $a_i, b_i \in [0, 1]$ with $a_i + b_i \in [0, 1]$, define a sequence of probability measures, $\boldsymbol{p}_i,\ i = 1, ..., n$ by setting

$$\boldsymbol{p}_i(A_1) = a_i$$

$$\boldsymbol{p}_i(A_2) = b_i$$

$$\boldsymbol{p}_i(A_3) = 1 - a_i - b_i$$

Since it is the case that $\mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) \in \mathbb{P}$ for all $(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \in \mathbb{P}^n$ and every $\mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) \in \mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, we have it that $\mathfrak{m}\big(\boldsymbol{p}_1(A), ..., \boldsymbol{p}_n(A)\big) = \boldsymbol{p}(A)$, for some $\boldsymbol{p} \in \mathbb{P}$ and all $A \in \mathscr{A}$. Now, by the additivity of probability measures, $\boldsymbol{p}(A_1 \cup A_2) = \boldsymbol{p}(A_1) + \boldsymbol{p}(A_2)$. Hence, $\mathfrak{m}(a_1 + b_1, ..., a_n + b_n) = \mathfrak{m}\big(a_1, ..., a_n\big) + \mathfrak{m}\big(b_1, ..., b_n\big)$. So, $\mathfrak{m}$ satisfies Cauchy's multivariable functional equation. For each $i = 1, ..., n$, define $\mathfrak{m}_i(a) = \mathfrak{m}(0, ..., a, ..., 0)$, where a occupies the i-th position of the vector $(0, ..., a, ..., 0)$. It is clear that $\mathfrak{m}_i(a + b) = \mathfrak{m}_i(a) + \mathfrak{m}_i(b)$ for all $a, b \in [0, 1]$ with $a + b \in [0, 1]$. Because $\mathfrak{m}$ is nonnegative, so is $\mathfrak{m}_i,\ i = 1, ..., n$. By Theorem 3 of (Aczél and Oser, 2006, p. 48), it follows that there exists a nonnegative constant α_i such that $\mathfrak{m}_i(a) = \alpha_i a$ for all $a \in [0,1]$. By the Cauchy equation we have

$$\begin{aligned} \mathfrak{m}(a_1, ..., a_n) &= \mathfrak{m}(a_1, 0, ..., 0) + \mathfrak{m}(0, a_2, ..., a_n) \\ &= \mathfrak{m}(a_1, 0, ..., 0) + \mathfrak{m}(0, a_2, 0, ..., 0) + ... + \mathfrak{m}(0, ..., 0, a_n) \end{aligned}$$

So we have $\mathfrak{m}(a_1, ..., a_n) = \mathfrak{m}_1(a_1) + ... + \mathfrak{m}_n(a_n) = \alpha_1 a_1 + ... + \alpha_n a_n$. And since $\mathfrak{m}(1, ..., 1) = 1$ (by consideration of the probability of Ω), it follows that $\sum_{i=1}^n \alpha_i = 1$. Thus, $\mathfrak{m}$ is a convex combination.

Now, we want to show that $conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \subseteq \mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Let $\boldsymbol{p}$ be an element of $conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. It is clear that there exists an $\mathfrak{m} \in \mathfrak{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ such that $\mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) = \boldsymbol{p}$. And since $\boldsymbol{p}$ is just a convex combination, there exist weights $\alpha_1, ..., \alpha_n \in [0,1]$ such that $\sum_{i=1}^n \alpha_i = 1$ and $\boldsymbol{p} = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i$. But for any other profile $(\boldsymbol{q}_1, ..., \boldsymbol{q}_n) \in \mathbb{P}^n$, taking any convex combination yields a probability measure. In particular, $\sum_{i=1}^n \alpha_i \boldsymbol{q}_i \in \mathbb{P}$. It follows that $\mathfrak{m} \in \bigcap_{\vec{\boldsymbol{q}} \in \mathbb{P}^n} \mathfrak{M}_n(\vec{\boldsymbol{q}})$. So, $\boldsymbol{p} = \mathfrak{m}\big(\boldsymbol{p}_1(\cdot), ..., \boldsymbol{p}_n(\cdot)\big) \in \mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, as desired.

The two inclusions above show that $\mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. Hence, $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = conv(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ is equivalent to $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{M}_n(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. $\square$

APPENDIX B

Appendix to Chapter 3: Proofs

Proof of Proposition 8

Proof. We follow through Wagner's proof for the precise case (2009, Theorem 3.3), adapting it for IP where necessary.

(⇒) Assume that $\mathcal{F}$ is externally Bayesian, i.e., for all profiles and any likelihood function, $\mathcal{F}^{\lambda}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}(\boldsymbol{p}_1^{\lambda}, ..., \boldsymbol{p}_n^{\lambda})$. We want to show that, for all partitions $\boldsymbol{E} = \{E_k\}$ of Ω and all profiles in $\mathbb{P}^n$,

$$\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \left\{ \frac{\sum_k b_k \boldsymbol{p}[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}(E_k)} : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \right\}$$

$$= \mathcal{F}\left(\frac{\sum_k b_k \boldsymbol{p}_1[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}_1(E_k)}, ..., \frac{\sum_k b_k \boldsymbol{p}_n[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}_n(E_k)} \right)$$

$$= \mathcal{F}(\boldsymbol{p}_{1J}^{\boldsymbol{E}}, ..., \boldsymbol{p}_{nJ}^{\boldsymbol{E}})$$

where the first and last equalities are definitional. Recall the definition of b_k: $b_k = \mathcal{B}(\boldsymbol{q}, \boldsymbol{p}; E_k : E_1) = \dfrac{\boldsymbol{q}(E_k)/\boldsymbol{q}(E_1)}{\boldsymbol{p}(E_k)/\boldsymbol{p}(E_1)}$, $k = 1, 2, ...$ Set $\lambda(\omega) = \sum_k b_k[\omega \in E_k]$. Wagner observes the following chain of equalities then obtains for $\boldsymbol{p}_i, i = 1, ..., n$ (2009, (3.10), p. 342):

$$(\star) \sum_{\omega \in \Omega} \lambda(\omega)\boldsymbol{p}_i(\omega) = \sum_{\omega \in \Omega} \boldsymbol{p}_i(\omega) \sum_k b_k[\omega \in E_k] = \sum_k b_k \sum_{\omega \in \Omega} \boldsymbol{p}_i(\omega)[\omega \in E_k] = \sum_k b_k \boldsymbol{p}_i(E_k)$$

Since each of the terms $b_k \boldsymbol{p}_i(E_k)$ is positive and $\sum_k b_k \boldsymbol{p}_i(E_k) < \infty$, λ is a likelihood function for $\boldsymbol{p}_i$, with $\boldsymbol{p}_i^{\lambda}$ a defined, updated pmf for $i = 1, ..., n$. Using $(\star)$, we can obtain

$$\mathcal{F}(\boldsymbol{p}_{1J}^{\boldsymbol{E}}, ..., \boldsymbol{p}_{nJ}^{\boldsymbol{E}}) = \mathcal{F}\left(\frac{\boldsymbol{p}_1 \lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}_1(\omega')\lambda(\omega')}, ..., \frac{\boldsymbol{p}_n \lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}_n(\omega')\lambda(\omega')} \right)$$

by substituting, for each $i = 1, ..., n$, $\lambda(\cdot)$ for $\sum_k b_k[\omega \in E_k]$ in the numerator and $\sum_{\omega' \in \Omega} \boldsymbol{p}_i(\omega')\lambda(\omega')$ for $\sum_k b_k \boldsymbol{p}_i(E_k)$ in the denominator. But by definition,

$$\mathcal{F}\left(\frac{\boldsymbol{p}_1\lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}_1(\omega')\lambda(\omega')}, ..., \frac{\boldsymbol{p}_n\lambda(\cdot)}{\sum_{\omega' \in \Omega} \boldsymbol{p}_n(\omega')\lambda(\omega')}\right) = \mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda)$$

and by assumption $\mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda) = \mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. By definition, $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \{\boldsymbol{p}^\lambda : \boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)\}$. But, for all $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, $\boldsymbol{p}^\lambda = \frac{\sum_k b_k \boldsymbol{p}[\cdot \in E_k]}{\sum_k b_k \boldsymbol{p}(E_k)}$. Hence, $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. So, $\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}(\boldsymbol{p}_{1J}^{\boldsymbol{E}}, ..., \boldsymbol{p}_{nJ}^{\boldsymbol{E}})$ follows from the assumption.

($\Leftarrow$) Suppose that $\mathcal{F}$ satisfies CJC_W and that λ is a likelihood function for $\boldsymbol{p}_i, i = 1, ..., n$. Let $(\omega_1, \omega_2, ...)$ be a list of all of those $\omega \in \Omega$ such that $\lambda(\omega) > 0$, and let $\boldsymbol{E} = \{E_1, E_2, ...\}$, where $E_i := \{\omega_i\}$. Setting $b_k = \frac{\lambda(\omega_k)}{\lambda(\omega_1)}$ for $k = 1, 2, ...$, it follows that $b_k > 0$ and that $b_1 = 1$. Since λ is a likelihood for $\boldsymbol{p}_i, i = 1, ..., n$, we have $\sum_k b_k \boldsymbol{p}_i(E_k) < \infty, i = 1, ..., n$, and that $(\boldsymbol{q}_1, ..., \boldsymbol{q}_n) \in \mathbb{P}^n$, where $\boldsymbol{q}_i(\omega) := \frac{\sum_k b_k \boldsymbol{p}_i(\omega)[\omega \in E_k]}{\sum_k b_k \boldsymbol{p}_i(E_k)}$. From CJC_W, it follows that 1) $0 < \sum_k b_k \boldsymbol{p}(E_k) < \infty$ for all $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, and that 2) $\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}(\boldsymbol{p}_{1J}^{\boldsymbol{E}}, ..., \boldsymbol{p}_{nJ}^{\boldsymbol{E}})$. 1) implies that $0 < \sum_{\omega \in \Omega} \lambda(\omega)\boldsymbol{p}(\omega) < \infty$ for all $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, and 2) implies that $\mathcal{F}^\lambda(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) = \mathcal{F}(\boldsymbol{p}_1^\lambda, ..., \boldsymbol{p}_n^\lambda)$ (since substituting the definition of b_k in terms of λ in $\frac{\sum_k b_k \boldsymbol{p}_i(\omega)[\omega \in E_k]}{\sum_k b_k \boldsymbol{p}_i(E_k)}$, the formula for obtaining the $\boldsymbol{q}_i$, reduces that formula to the formula for updating on that λ).

□

Proof of Proposition 11

Proof. We provide a case in which convex IP pooling and Jeffrey conditionalization *as standardly construed* do not commute. Let $\boldsymbol{q}_i$ come from $\boldsymbol{p}_i$ by Jeffrey conditionalization, and let $\boldsymbol{q}$ be a common posterior distribution over partition $\boldsymbol{E}$ for $\boldsymbol{p}_i$, $i = 1, ..., n$. Let $\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ come from $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ by Jeffrey conditionalizing each $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ using $\boldsymbol{q}$, the common posterior distribution over $\boldsymbol{E}$. We offer a counterexample to commutativity in which $\mathcal{F}_J^{\boldsymbol{E}}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n) \neq \mathcal{F}(\boldsymbol{q}_1, ..., \boldsymbol{q}_n)$.

Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, and consider the following two pmfs.

TABLE 1. Priors

	ω_1	ω_2	ω_3	ω_4
$\boldsymbol{p}_1$	1/4	1/4	1/4	1/4
$\boldsymbol{p}_2$	1/8	1/2	1/4	1/8

Let $\boldsymbol{E} = \{E_1, E_2\}$ with $E_1 = \{\omega_1, \omega_2\}$ and $E_2 = \{\omega_3, \omega_4\}$ be a partition of Ω. Jeffrey updating both pmfs using $\boldsymbol{q}$, where $\boldsymbol{q}(E_1) = 2/3$ and $\boldsymbol{q}(E_2) = 1/3$, we obtain the following posteriors.

TABLE 2. Posteriors

	ω_1	ω_2	ω_3	ω_4
$\boldsymbol{q}_1$	1/3	1/3	1/6	1/6
$\boldsymbol{q}_2$	2/15	8/15	2/9	1/9

Consider the $.50 - .50$ mixture of $\boldsymbol{p}_1$ and $\boldsymbol{p}_2$, $\boldsymbol{p}^\star = 0.5\boldsymbol{p}_1 + 0.5\boldsymbol{p}_2$. It is clear that $\boldsymbol{p}^\star \in \mathcal{F}(\boldsymbol{p}_1, \boldsymbol{p}_2)$. Jeffrey conditionalizing $\boldsymbol{p}^\star$ with $\boldsymbol{q}$ gives us $\boldsymbol{q}^\star$. In particular, $\boldsymbol{q}^\star(\omega_1) = 2/9$ and $\boldsymbol{q}^\star(\omega_3) = 4/21$. It is clear that $\boldsymbol{q}^\star \in \mathcal{F}^{J}_{\boldsymbol{E}}(\boldsymbol{p}_1, \boldsymbol{p}_2)$. Any $\boldsymbol{q}_\star \in \mathcal{F}(\boldsymbol{q}_1, \boldsymbol{q}_2)$ is of the form $\boldsymbol{q}_\star = \alpha\boldsymbol{q}_1 + (1-\alpha)\boldsymbol{q}_2$ for $\alpha \in [0, 1]$.

Suppose that $\mathcal{F}^{\boldsymbol{E}}_{J}(\boldsymbol{p}_1, \boldsymbol{p}_2) = \mathcal{F}(\boldsymbol{q}_1, \boldsymbol{q}_2)$. Then, there is a $\boldsymbol{q}_\star \in \mathcal{F}(\boldsymbol{q}_1, \boldsymbol{q}_2)$ such that $\boldsymbol{q}^\star = \boldsymbol{q}_\star$. In particular, $\boldsymbol{q}_\star(\omega_1) = 2/9$ and $\boldsymbol{q}_\star(\omega_3) = 4/21$. Letting $\boldsymbol{q}_\star(\omega_1) = 2/9$, we can compute α.

$$2/9 = \boldsymbol{q}_\star(\omega_1) = \alpha\boldsymbol{q}_1(\omega_1) + (1-\alpha)\boldsymbol{q}_2(\omega_1) = \alpha 1/3 + (1-\alpha)2/15$$

Solving, we get $\alpha = 4/9$. However, we are supposed to have $\boldsymbol{q}_\star(\omega_3) = 4/21$. For $\alpha = 4/9$, that is not the case.

$$\boldsymbol{q}_\star(\omega_3) = \alpha\boldsymbol{q}_1(\omega_3) + (1-\alpha)\boldsymbol{q}_2(\omega_3) = 4/9(1/6) + 5/9(2/9) = 16/81 > 4/21 = \boldsymbol{q}^\star(\omega_3)$$

It follows that $\mathcal{F}^{\boldsymbol{E}}_{J}(\boldsymbol{p}_1, \boldsymbol{p}_2) \neq \mathcal{F}(\boldsymbol{q}_1, \boldsymbol{q}_2)$.

□

Proof of Proposition 12

PROOF. We want to show that $\mathcal{F}(\boldsymbol{q}_1, ..., \boldsymbol{q}_n) = \mathcal{F}_I^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, where $\boldsymbol{q}_i$ comes from $\boldsymbol{p}_i$ by general imaging on E, and $\mathcal{F}_I^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ comes from $\mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ by general imaging each $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ on E. Again, we show both inclusions. In the proofs, we appeal to the fact any element of a convex set is some convex combination of the generating, extreme points: For any $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n), \boldsymbol{p} = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i$, where $\alpha_i \geqslant 0$ for $i = 1, ..., n$, and $\sum_{i=1}^n \alpha_i = 1$ (see, e.g., Redacted, Redacted, Lemma 1).

Let $\boldsymbol{q} \in \mathcal{F}(\boldsymbol{q}_1, ..., \boldsymbol{q}_n)$. So, $\boldsymbol{q} = \sum_{i=1}^n \alpha_i \boldsymbol{q}_i$. Since $\boldsymbol{q}$ is a linear pool of $\boldsymbol{q}_i$ for $i = 1, ..., n$, by Gärdenfors' result, Theorem 10, $\boldsymbol{q}$ is also the result of imaging $\boldsymbol{p} = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i$ on E, because linear pooling and general imaging commute. Since $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$, it follows that $\boldsymbol{q} \in \mathcal{F}_I^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$.

For the other direction, assume that $\boldsymbol{q} \in \mathcal{F}_I^E(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$. So, $\boldsymbol{q}$ is the result of general imaging some $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n)$ on E. For any $\boldsymbol{p} \in \mathcal{F}(\boldsymbol{p}_1, ..., \boldsymbol{p}_n), \boldsymbol{p} = \sum_{i=1}^n \alpha_i \boldsymbol{p}_i$. By Gärdenfors' result, $\boldsymbol{q} = \sum_{i=1}^n \alpha_i \boldsymbol{q}_i$, where the $\boldsymbol{q}_i$ come from the $\boldsymbol{p}_i$ by general imaging on E, because general imaging and linear pooling commute. But then it follows that $\boldsymbol{q} \in \mathcal{F}(\boldsymbol{q}_1, ..., \boldsymbol{q}_n)$.

□

APPENDIX C

Appendix to Chapter 4: Proofs

Proof of Proposition 13

PROOF. Let $\mathbb{P}^1, ..., \mathbb{P}^k$ be the extreme points that generate C. We show that for almost every $\omega \in \Omega$,

$$\sup_{\mathbb{P},\mathbb{Q}\in C} d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \to 0$$

as $n \to \infty$. Note that all of our "almost surely" qualifications hold for every $\mathbb{P} \in C$ by the mutual absolute continuity assumption.

Let $\mathbb{P}, \mathbb{Q} \in C$ be arbitrary. Following Huttegger (2015a, Appendix), we note that for all $\omega \in \Omega$ the Jeffrey conditioning equation (6) can be written as

$$\mathbb{P}_n[\cdot](\omega) = \int \mathbb{P}[\cdot \mid \mathfrak{F}_n] dP_n(\omega). \tag{7}$$

Using (7) and our assumption that $P_n(\omega) = P_n^1(\omega) = Q_n(\omega)$ for almost every ω, we have

$$d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \leqslant \int d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) dP_n^1(\omega) = \int d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \frac{dP_n^1}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1 \tag{8}$$

for almost all ω. The function $dP_n^1(\omega)/d\mathbb{P}^1$ is the Radon-Nikodym derivative of P_n^1 with respect to $\mathbb{P}^1$, which is guaranteed to exist for almost every ω because $P_n^1(\omega)$ is absolutely continuous with respect to $\mathbb{P}^1$ for almost every ω. Note that

$$d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \leqslant \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \text{ a.s.}$$

because $\mathbb{P}_{\mathfrak{F}_n}$ and $\mathbb{Q}_{\mathfrak{F}_n}$ are (almost surely) convex combinations of $\mathbb{P}^1_{\mathfrak{F}_n}, ..., \mathbb{P}^k_{\mathfrak{F}_n}$ (Schervish and Seidenfeld, 1990). Therefore, for almost every ω,

$$\int d(\mathbb{P}_{\mathfrak{F}_n}, \mathbb{Q}_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1 \leqslant \int \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1.$$

From this inequality and (8) we have

$$d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \leqslant \int \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1 \tag{9}$$

for almost every ω. Since (9) holds for all $\mathbb{P}, \mathbb{Q} \in C$, it holds upon taking a supremum over $\mathbb{P}, \mathbb{Q} \in C$. Hence, for almost all ω we have

$$\sup_{\mathbb{P}, \mathbb{Q} \in C} d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \leqslant \int \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1. \tag{10}$$

We conclude the proof with arguments similar to Huttegger's. We refer the reader to his Appendix for more details. First, we note that (M') implies that the sequence $\{dP^1_n(\omega)/d\mathbb{P}^1\}$ is a nonnegative martingale with respect to $\mathbb{P}^1$ for almost every ω, and hence converges almost surely to a finite limit for almost every ω. Theorem 12 implies that $\max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \to 0$ almost surely. Therefore, for almost every ω,

$$\max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) \to 0$$

almost surely as $n \to \infty$. Using this fact we have

$$\lim_{n\to\infty} \int \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1 = \int \lim_{n\to\infty} \max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) \frac{dP^1_n}{d\mathbb{P}^1}(\omega) d\mathbb{P}^1 = 0 \tag{11}$$

for almost every ω. The passage of the limit under the integral is justified by the fact that $dP^1_n(\omega)/d\mathbb{P}^1$ is uniformly integrable for almost every ω, which in turn implies that $\max_{i,j} d(\mathbb{P}^i_{\mathfrak{F}_n}, \mathbb{P}^j_{\mathfrak{F}_n}) dP^1_n(\omega)/d\mathbb{P}^1$ is uniformly integrable for almost every ω as it is dominated by $dP^1_n(\omega)/d\mathbb{P}^1$. That $dP^1(\omega)_n/d\mathbb{P}^1$

is uniformly integrable for almost every ω is equivalent to our assumption that $P_n^1(\omega)$ is uniformly absolute continuous with respect to $\mathbb{P}^1$ for almost every ω. See Huttegger (2015a, Proof of Theorem 9.2) for more on this point.

Finally, (10) and (11) imply

$$\sup_{\mathbb{P},\mathbb{Q}\in C} d(\mathbb{P}_n(\omega), \mathbb{Q}_n(\omega)) \to 0$$

as $n \to \infty$ for almost every ω, as desired. □

Jeffrey Conditioning Does Not Preserve Convexity

PROPOSITION 27. *Let C be as in Proposition 13. The result of Jeffrey conditioning all elements of C on a common posterior need not be a convex set.*

PROOF. We sketch a counterexample. Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, and consider the following two probabilities.

TABLE 1. Priors

	ω_1	ω_2	ω_3	ω_4
$\mathbb{P}$	1/4	1/4	1/4	1/4
$\mathbb{Q}$	1/8	1/2	1/4	1/8

Set $C = conv\{\mathbb{P}, \mathbb{Q}\}$. Let $\mathfrak{E} = \{E_1, E_2\}$ with $E_1 = \{\omega_1, \omega_2\}$ and $E_2 = \{\omega_3, \omega_4\}$ be a partition of Ω. Jeffrey updating both $\mathbb{P}$ and $\mathbb{Q}$ using P, where $P(E_1) = 2/3$ and $P(E_2) = 1/3$, we obtain the following posteriors.

Table 2. Posteriors

	ω_1	ω_2	ω_3	ω_4
$\mathbb{P}_1$	1/3	1/3	1/6	1/6
$\mathbb{Q}_1$	2/15	8/15	2/9	1/9

Let $C^P_{\mathfrak{E}}$ be the result of Jeffrey updating each element of C on the common posterior P on partition $\mathfrak{E}$. To establish the claim, it suffices to find some $\alpha \in [0,1]$ such that $\alpha\mathbb{R}_1 + (1-\alpha)\mathbb{R}'_1 \notin C^P_{\mathfrak{E}}$ for some $\mathbb{R}, \mathbb{R}' \in C$ (since it is clear that $\mathbb{R}_1, \mathbb{R}'_1 \in C^P_{\mathfrak{E}}$). Let $\alpha = 4/9$ and consider $\alpha\mathbb{P}_1 + (1-\alpha)\mathbb{Q}_1$. Suppose for *reductio* that $C^P_{\mathfrak{E}}$ is convex. Then, $\alpha\mathbb{P}_1 + (1-\alpha)\mathbb{Q}_1 \in C^P_{\mathfrak{E}}$. This implies that there is some $\beta \in [0,1]$ and $\mathbb{R}^* \in C$ such that $\beta\mathbb{P} + (1-\beta)\mathbb{Q} = \mathbb{R}^*$ and $\mathbb{R}^*_1 = \alpha\mathbb{P}_1 + (1-\alpha)\mathbb{Q}_1$.

Clearly, $\alpha\mathbb{P}_1(\omega_3) + (1-\alpha)\mathbb{Q}_1(\omega_3) = 16/81$. By the definition of Jeffrey conditioning and our assumptions, we have

$$\frac{16}{81} = \sum_{j=1,2} P(E_j)\mathbb{R}^*(\omega_3|E_j) = \left(\frac{2}{3}\right)\frac{\mathbb{R}^*(\omega_3 \cap E_1)}{\mathbb{R}^*(E_1)} + \left(\frac{1}{3}\right)\frac{\mathbb{R}^*(\omega_3 \cap E_2)}{\mathbb{R}^*(E_2)}.$$

Since $\omega_3 \cap E_1 = \varnothing$, the left summand of the right-hand side is 0. By algebra and the definition of $\mathbb{R}^*$, we obtain

$$\frac{16}{81} = \left(\frac{1}{3}\right)\frac{\beta\mathbb{P}(\omega_3) + (1-\beta)\mathbb{Q}(\omega_3)}{\beta\mathbb{P}(E_2) + (1-\beta)\mathbb{Q}(E_2)}.$$

Substituting values from Table 1 and solving for β, we have $\beta = 3/8$. However, $\alpha\mathbb{P}_1(\omega_1) + (1-\alpha)\mathbb{Q}_1(\omega_1) = 2/9$. This implies that $\mathbb{R}^*_1(\omega_1) = 2/9$. But for P and $\beta = 3/8$, this is not the case. Again using the definition of Jeffrey conditioning and our assumptions, it can be verified that $\mathbb{R}^*_1(\omega_1) = 22/111 < 2/9$. It follows that there do not exist a $\beta \in [0,1]$ and a $\mathbb{R}^* \in C$ that meet our stated requirements above. □

APPENDIX D

Appendix to Chapter 5: Proofs

Proof of Proposition 15

PROOF. (cf. Hansson, 1999, p. 135) Let $\mathcal{L}$ be based on two atomic sentences, p and q. $\mathcal{L}$ contains sixteen non-equivalent sentences. Let $K_1 = Cn(\{p \wedge q\})$. Then, we can list the logically distinct elements of K_1: $\{p \wedge q, p, p \leftrightarrow q, q, p \to q, p \vee q, q \to p, \top\}$. Let $K_2 = Cn(\{p \vee q\})$. Then, $K_1 \cap K_2 = K_2$. Consider the remainder sets of K_1.

$$\begin{aligned}
K_1 \perp p \wedge q &= \{Cn(\{p\}), Cn(\{q\}), Cn(\{p \leftrightarrow q\})\} \\
K_1 \perp p &= \{Cn(\{q\}), Cn(\{p \leftrightarrow q\})\} \\
K_1 \perp p \leftrightarrow q &= \{Cn(\{p\}), Cn(\{q\})\} \\
K_1 \perp q &= \{Cn(\{p\}), Cn(\{p \leftrightarrow q\})\} \\
K_1 \perp p \to q &= \{Cn(\{p\})\} \\
K_1 \perp p \vee q &= \{Cn(\{p \leftrightarrow q\})\} \\
K_1 \perp q \to p &= \{Cn(\{q\})\} \\
K_1 \perp \top &= \varnothing
\end{aligned}$$

Now let γ be such that

$$\begin{aligned}
\gamma(K_1 \perp p \wedge q) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma(K_1 \perp p) &= \{Cn(\{q\}), Cn(\{p \leftrightarrow q\})\} \\
\gamma(K_1 \perp p \leftrightarrow q) &= \{Cn(\{q\})\} \\
\gamma(K_1 \perp q) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma(K_1 \perp p \rightarrow q) &= \{Cn(\{p\})\} \\
\gamma(K_1 \perp p \vee q) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma(K_1 \perp q \rightarrow p) &= \{Cn(\{q\})\} \\
\gamma(K_1 \perp \top) &= \{Cn(\{p \wedge q\})\}
\end{aligned}$$

With the exception of $K_1 \dot{-} p$, the partial meet contractions are given by the sole members of the selection sets. $K_1 \dot{-} q = \bigcap \gamma(K_1 \perp q)$, for instance, is $Cn(\{p \leftrightarrow q\})$. For $K_1 \dot{-} p$, we have $\bigcap\{Cn(\{q\}), Cn(\{p \leftrightarrow q\})\} = Cn(\{p \rightarrow q\})$. So, not only do we have $K_1 \dot{-} p \vee q \neq K_1 \cap K_2$, but $K_1 \dot{-} \alpha \neq K_1 \cap K_2$ for any α. In this example, $\dot{-}$ is *not* a maxichoice contraction. Another example sufficient to illustrate the proposition could be constructed by selecting a single element of $K_1 \perp p$, say $p \leftrightarrow q$, converting $\dot{-}$ into a maxichoice contraction (which would also be a transitively relational partial meet contraction), but the present example illustrates that the limitations for shifting to consensus apply to partial meet contractions more generally. □

Proof of Proposition 16

Proof. Let K_1 be a maximal consistent subset of $\mathcal{L}$, i.e., for any $\alpha \in \mathcal{L}$, either $\alpha \in K_1$ or $\neg\alpha \in K_1$ (but not both). Let $K_2 = Cn(\varnothing)$. We show that the *full meet* contraction of K_1 by any sentence α—i.e., $\gamma(K_1 \perp \alpha) = K_1 \perp \alpha$, so $\bigcap \gamma(K_1 \perp \alpha) = \bigcap K_1 \perp \alpha$—is not identical to $Cn(\varnothing)$. Specifically, we show that any full meet is a superset of $Cn(\varnothing)$. And since any sentence in the full meet is also in any partial meet, it follows that no partial meet contraction of K_1 is identical to $Cn(\varnothing) = K_1 \cap K_2$.

If $\dot{-}$ is a full meet contraction function and $\alpha \in K_1$, then $K_1 \dot{-} \alpha = K_1 \cap Cn(\{\neg\alpha\})$ (e.g., Gärdenfors, 1988, Lemma 4.9; Hansson, 1999, Observation 2.12). α contains instances of only finitely many atomic sentences, where $\{p_i : i \in I\}$ is the set of atomic sentences of $\mathcal{L}$. Let p_i be an atomic sentence in K_1 that does *not* occur in α. Clearly, $\neg\alpha \vee p_i \in K_1 \cap Cn(\{\neg\alpha\})$. If $\alpha \in Cn(\varnothing)$, then $K_1 \cap Cn(\{\neg\alpha\}) = K_1 \neq Cn(\varnothing)$. So, let α be contingent. Since p_i is atomic, $p_i \notin Cn(\varnothing)$. But since all of the atomic sentences are independent and p_i does not occur in α, $\neg\alpha \vee p_i \notin Cn(\varnothing)$ either (consider the truth table). So $K_1 \cap K_2 = Cn(\varnothing) \subset K_1 \dot{-} \alpha$ for any sentence α. □

Unanimous Consensus for Package Contraction

Motivated by the limitations of "singleton" contraction, Fuhrmann and Hansson study *multiple* contractions, operations simultaneously contracting a *set* of sentences (1994). They identify two primary types of multiple contractions. *Package* contractions remove all of a set of sentences from the belief set; *choice* contractions remove at least one among a set of sentences. As they argue, multiple contraction should not be thought of in terms of sequential contraction. In sequential contractions, the order of contractions makes a difference to the resulting belief state. So a privileged sequence would have to be identified. Such a sequence finds no analogue in multiple contraction. This is to say nothing of the infamous problems confronting AGM regarding iterated revisions.

The setup for partial meet package contraction mirrors that of AGM partial meet contraction. We make use of a multiple conclusion relation, $\vdash$, familiar from sequent calculus. For sets of sentences, A, B, we put $A \vdash B$ when $B \cap Cn(A) \neq \varnothing$. The postulates below make use of a notion of equivalence-according-to-K. $A \equiv_K B$ holds *iff* $\forall C \subseteq K : C \vdash A \Leftrightarrow C \vdash B$.

Fuhrmann and Hansson's postulates are as follows.

(−1) $K - [A] \subseteq K$

(−2) If $\varnothing \nvdash A$, then $A \cap (K - [A]) = \varnothing$

(−3) If $A \equiv_K B$, then $K - [A] = K - [B]$

(−4) If $\alpha \in K \backslash K - [A]$, then there is some K' such that $K - [A] \subseteq K' \subseteq K$ and $K' \nvdash A$ and $K', \alpha \vdash A$

Definition 7. The set of all package A-remainders of K is given by

$$K \bot A = \{K' \subseteq K : (i) K' \nvdash A, (ii) \forall K'' : K' \subset K'' \subseteq K \Rightarrow K'' \vdash A\}$$

A package partial meet contraction takes the intersection of some selection of elements in $K \bot A$

$$K - [A] = \bigcap \gamma(K \bot A)$$

Fuhrmann and Hansson's postulates characterize package partial meet contraction.

Theorem 20. *(Fuhrmann and Hansson, 1994, Thm. 9) An operation $-$ is a $\bot$ partial meet contraction iff $-$ satisfies* $(-1)-(-4)$ *for all $K \in \mathbb{K}$ and all sets $A, B \in \mathscr{P}(\mathcal{L})$.*

Proposition 28. *Let $-$ be a package partial meet contraction operator. If $A = K_1 \backslash K_2$, then $K_1 - [A] = K_1 \cap K_2$.*

Proof. The proof proceeds by showing that $K_1 \cap K_2$ is the only A-remainder of K_1, i.e., $K_1 \bot A = \{K_1 \cap K_2\}$. First, then, we show that $K_1 \cap K_2 \in K_1 \bot A$. It is clear that both $(K_1 \cap K_2) \subseteq K_1$ and that $K_1 \cap K_2 \nvdash A$. We need only show that for any $K' \in \mathbb{K}$, if $(K_1 \cap K_2) \subset K' \subseteq K_1$, then $K' \vdash A$. Suppose the antecedent holds. Then, there is some $\alpha \in K'$ such that $\alpha \notin K_1 \cap K_2$. It follows that $\alpha \in K_1 \backslash K_2 = A$. Hence, $K' \vdash A$. So, $K_1 \cap K_2 \in K_1 \bot A$.

Assume for *reductio* that there is a $K' \in K_1 \bot A$ such that $K' \neq K_1 \cap K_2$. So either there is some $\alpha \in K' \backslash (K_1 \cap K_2)$, or else there is some $\alpha \in (K_1 \cap K_2) \backslash K'$. Suppose the former. Since $K' \subseteq K_1$, α must be in A. So, $K' \vdash A$, contrary to our assumption that K' is an A-remainder of K_1. Suppose the latter. Then, the only possibility is that $K' \subset (K_1 \cap K_2) \subseteq K_1$ (since $K' \nsubseteq (K_1 \cap K_2)$ entails that there is an $\beta \in K' \backslash (K_1 \cap K_2)$, and, again, $\beta \in K' \backslash (K_1 \cap K_2)$ entails $K' \vdash A$). But clearly, $K_1 \cap K_2 \nvdash A$, again, contrary to our assumption that K' is an A-remainder of K_1. $\square$

APPENDIX E

Appendix to Chapter 6: Proofs

Proof of Proposition 18

Proof. Let $\mathcal{L}$ be the language containing only the truth-functional combinations of two atomic sentences, p and q. $\mathcal{L}$, then, contains sixteen non-equivalent sentences. Let $K_1 = Cn(\{p \wedge q\})$. The set of logically distinct elements of K_1 is

$$\{p \wedge q, p, p \leftrightarrow q, q, p \rightarrow q, p \vee q, q \rightarrow p, \top\}$$

We can list the saturatable subsets of K_1.

$$\begin{aligned}
S(K_1, p \wedge q) &= \{Cn(\{p\}), Cn(\{q\}), Cn(\{p \leftrightarrow q\})\} \\
S(K_1, p) &= \{Cn(\{q\}), Cn(\{q \rightarrow p\}), Cn(\{p \vee q\}), Cn(\{p \leftrightarrow q\})\} \\
S(K_1, p \leftrightarrow q) &= \{Cn(\{p\}), Cn(\{q\}), Cn(\{q \rightarrow p\}), Cn(\{p \rightarrow q\})\} \\
S(K_1, q) &= \{Cn(\{p\}), Cn(\{p \rightarrow q\}), Cn(\{\neg p \rightarrow q\}), Cn(\{p \leftrightarrow q\})\} \\
S(K_1, p \rightarrow q) &= \{Cn(\{p\}), Cn(\varnothing)\} \\
S(K_1, p \vee q) &= \{Cn(\{p \leftrightarrow q\}), Cn(\varnothing)\} \\
S(K_1, q \rightarrow p) &= \{Cn(\{q\}), Cn(\varnothing)\} \\
S(K_1, \top) &= \varnothing
\end{aligned}$$

Now let γ be such that

$$\begin{aligned}
\gamma\big(S(K_1, p \wedge q)\big) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma\big(S(K_1, p)\big) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma\big(S(K_1, p \leftrightarrow q)\big) &= \{Cn(\{q \rightarrow p\})\} \\
\gamma\big(S(K_1, q)\big) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma\big(S(K_1, p \rightarrow q)\big) &= \{Cn(\{\varnothing\})\} \\
\gamma\big(S(K_1, p \vee q)\big) &= \{Cn(\{p \leftrightarrow q\})\} \\
\gamma\big(S(K_1, q \rightarrow p)\big) &= \{Cn(\{\varnothing\})\} \\
\gamma\big(S(K_1, \top)\big) &= \{Cn(\{p \wedge q\})\}
\end{aligned}$$

Since γ returns a singleton in every case, the meets of the selection sets above are trivial. For example, $\bigcap \gamma\big(S(K_1, p \wedge q)\big) = Cn(\{p \leftrightarrow q\})$. If $K_2 = Cn(\{p \vee q\})$, it is clear that for no sentence, α, do we have $K_1 \dot{-} \alpha = K_1 \cap K_2 = Cn(\{p \vee q\})$. □

Proof of Proposition 19

Proof. It suffices to observe that the Levi-contraction for K_1 in the proof of Proposition 18 is a severe withdrawal function based on the entrenchment ordering $\leqslant$ with tautologies as maximal elements, non-beliefs as minimal elements, and elements of K_1 (weakly) ordered as follows:

$$p \vee q \sim p \wedge q \sim p \sim q < p \leftrightarrow q \sim p \rightarrow q < q \rightarrow p$$

The partial meet contractions of Proposition 18 can be recovered from this entrenchment ordering *via* Definition 5. So, since $\leqslant$ is an epistemic entrenchment ordering (satisfying $(E1) - (E5)$), by Theorem 16, $\dot{-}$ satisfies the axioms that characterize mild contraction. If, again, $K_2 = Cn(\{p \vee q\})$, then $K_1 \cap K_2 = K_2 = Cn(\{p \vee q\})$. Yet, $K_1 \dot{-} \alpha \neq Cn(\{p \vee q\})$ for any $\alpha \in \mathcal{L}$. □

Proof of Proposition 20

Proof. Recall the relevant postulates:

($\dot{-}$3) If $\alpha \notin K$, then $K \subseteq K \dot{-} \alpha$

($\dot{-}$4) If $\alpha \in K \dot{-} \alpha$, then $\alpha \in Cn(\varnothing)$

($\dot{-}$9) If $\alpha \in Cn(\varnothing)$, then $K \subseteq K \dot{-} \alpha$

To see that ($\dot{-}$3) is not generally satisfied, suppose that $\alpha \notin K_1$ but that $\beta \in K_1 \backslash K_2$. Then, $K_1 \nsubseteq K_1 \dot{-}_{C_{K_2}} \alpha$ since $\beta \notin K_1 \cap K_2$. For a violation of ($\dot{-}$4), pick $\alpha \notin Cn(\varnothing)$. Suppose that $\alpha \in K_1 \cap K_2$. Then, $\alpha \in K_1 \dot{-}_{C_{K_2}} \alpha$. So ($\dot{-}$4) is violated. For ($\dot{-}$9), pick $\alpha \in Cn(\varnothing)$, but suppose $K_1 \backslash K_2 \neq \varnothing$. Then, $K_1 \dot{-}_{CK_2} \alpha = K_1 \cap K_2 \subset K_1$. Hence, $K_1 \nsubseteq K_1 \dot{-}_{C_{K_2}} \alpha$. $\square$

APPENDIX F

Appendix to Chapter 7: Proofs

Proof of Proposition 21

Proof. Let $(x,y) \in \mathcal{R}$. Then, $(x,y) \in \bigcap_{i\in N} P_i$ or $(x,y) \in \bigcap_{i\in N} I_i$. Either way, $(x,y) \in \bigcap_{i\in N} R_i$. □

Proof of Proposition 22

Proof. Let $\Delta = \{(x,x) : x \in X\}$ be the diagonal of X. Since $\Delta \subseteq I_i$ for all $i \in N$, it follows that $\Delta \subseteq \mathcal{R}$, so $\mathcal{R}$ is reflexive. For transitivity, assume that $x\mathcal{R}y$ and $y\mathcal{R}z$. There are four cases to consider to establish that $x\mathcal{R}z$.

(1) xP_iy and yP_iz for all $i \in N$, i.e., $(x,y),(y,z) \in \bigcap_{i\in N} P_i$. Since $\bigcap_{i\in N} P_i$ is transitive (Proposition 29), it follows that $(x,z) \in \bigcap_{i\in N} P_i$, i.e., xP_iz for all $i \in N$. Hence, $x\mathcal{R}z$.

(2) xP_iy and yI_iz for all $i \in N$. Suppose for *reductio* that zR_ix for some $i \in N$. By transitivity, we have that yR_ix, contradicting the assumption that xP_iy. So, xP_iz for all $i \in N$. Hence, $x\mathcal{R}z$.

(3) xI_iy and yP_iz for all $i \in N$. The proof is analogous to that for case (2).

(4) xI_iy and yI_iz for all $i \in N$. It follows from the transitivity of the R_i that xI_iz for all $i \in N$. Hence, $x\mathcal{R}z$.

□

Proof of Proposition 24

Proof. First, we want to show that for any $i \in N$, $R_i \in Ext(\mathcal{R})$. So, we need to show $\mathcal{R} \subseteq R_i$ and $\mathcal{R}\backslash\mathcal{R}^{-1} \subseteq R_i\backslash R_i^{-1}$. Let $(x,y) \in \mathcal{R} = \bigcap_{i\in N} P_i \cup \bigcap_{i\in N} I_i$. So, $(x,y) \in \bigcap_{i\in N} P_i$ or $(x,y) \in \bigcap_{i\in N} I_i$.

Either way, it follows $(x, y) \in R_i$. So, $\mathcal{R} \subseteq R_i$. And since $\mathcal{R} \backslash \mathcal{R}^{-1} = \bigcap_{i \in N} P_i$, it follows that $\mathcal{R} \backslash \mathcal{R}^{-1} \subseteq R_i \backslash R_i^{-1}$.

To establish the second claim of the proposition, we can use Example 1. Let $n = 2$ and $X = \{x, y, z\}$. Let R_1 be given by xP_1yP_1z and R_2 by xP_2yI_2z. $(y, z) \in \bigcap_{i \in N} R_i$, but $(z, y) \notin \bigcap_{i \in N} R_i$. Since $\bigcap_{i \in N} R_i \backslash \bigcap_{i \in N} R_i^{-1} \subseteq R \backslash R^{-1}$ for all $R \in Ext(\bigcap_{i \in N} R_i)$, it follows that $(y, z) \in R \backslash R^{-1}$. Because $(y, z) \notin R_2 \backslash R_2^{-1}$, $R_2 \notin Ext(\bigcap_{i \in N} R_i)$. □

Extensions of Categorical Strict Preference

PROPOSITION 29. *$\bigcap_{i \in N} P_i$ is a transitive relation.*

PROOF. Let $(x, y), (y, z) \in \bigcap_{i \in N} P_i$. That is, for all $i \in N$, xR_iy and $\neg yR_ix$ and yR_iz and $\neg zR_iy$. Since for all $i \in N$ R_i is a weak order, we have xR_iz by transitivity. zR_ix would imply zR_iy (again by transitivity), which contradicts the assumption that $\neg zR_iy$. So, xP_iz for all $i \in N$, i.e., $(x, z) \in \bigcap_{i \in N} P_i$. □

PROPOSITION 30. *Every transitive relation has a weak order extension.*

PROOF. Let $R \subseteq X \times X$ be a transitive relation. Let $\Delta = \{(x, x) : x \in X\}$ be the diagonal of X. Consider the relation $R' = R \cup \Delta$. R' is reflexive and transitive, i.e., a quasiorder. By Theorem 19, it follows that R' has a weak order extension that is also, then, an extension of R. □

To generate the elements of $Ext(\bigcap_{i \in N} P_i)$, begin by taking the union $\bigcap_{i \in N} P_i \cup \Delta$, obtaining a quasiorder.

Recall our generalization of the Pareto Unanimity to the context of set-valued SWFs. Where $P = R \backslash R^{-1}$,

Pareto Unanimity (P). For all $x, y \in X$ and all profiles in D, if xP_iy for $i = 1, ..., n$, then xPy for all $R \in \mathcal{F}(R_1, ..., R_n)$.

In the sense of the following propositions, $Ext(\bigcap_{i\in N} P_i)$ is the upper bound of $\mathcal{F}(R_1, ..., R_n)$ given (**P**) is imposed.

PROPOSITION 31. *If $\mathcal{F}$ satisfies **P**, then $\mathcal{F}(R_1, ..., R_n) \subseteq Ext(\bigcap_{i\in N} P_i)$ for any profile $(R_1, ..., R_n) \in D$.*

PROOF. Suppose that $\mathcal{F}$ satisfies **P**. Let $R \in \mathcal{F}(R_1, ..., R_n)$. We need to show that R is an extension of $\bigcap_{i\in N} P_i$. By **P**, $\bigcap_{i\in N} P_i \subseteq R$ for all $R \in \mathcal{F}(R_1, ..., R_n)$. Since $\bigcap_{i\in N} P_i = \bigcap_{i\in N} P_i \backslash \bigcap_{i\in N} P_i^{-1}$ and, by **P**, $\bigcap_{i\in N} P_i \subseteq R\backslash R^{-1}$ for all $R \in \mathcal{F}(R_1, ..., R_n)$, it follows that R is an extension of $\bigcap_{i\in N} P_i$. □

Levi's Categorical Preference

Isaac Levi argues for a slightly different notion of categorical preference (e.g., Levi, 2008). He seems not to understand weak preference as a derivative notion of strict preference and indifference. Accordingly, that y is weakly preferred to x in Example 1 is something to be preserved by categorical preference rather than an artifact to be dismissed because there is consensus neither that y is strictly preferred to x nor that y and x are indifferent. So, in addition to categorical strict preference and categorical indifference, Levi insists that all permissible orderings respect categorical weak preference as well. For any profile $(R_1, ..., R_n) \in \mathfrak{R}^n$, the largest set of permissible orderings according to Levi is given by L:

$$L = \{R \in \mathfrak{R} : (i) \bigcap_{i\in N} R_i \subseteq R, \text{ and } (ii) \bigcap_{i\in N} P_i \subseteq R\backslash R^{-1}\}$$

PROPOSITION 32. *For all $(R_1, ..., R_n) \in \mathfrak{R}^n$, $\{R_i : i \in N\} \subseteq L \subseteq Ext(\mathcal{R})$.*

PROOF. Let $(R_1, ..., R_n) \in \mathfrak{R}^n$. Since, for any $j \in N$, $\bigcap_{i\in N} R_i \subseteq R_j$ and $\bigcap_{i\in N} P_i \subseteq R_j\backslash R_j^{-1}$, the first inclusion holds. Condition (ii) is the same in the definitions of L and $Ext(\mathcal{R})$, so we need to only consider (i). Since $\mathcal{R} \subseteq \bigcap_{i \in N} R_i$ (Proposition 21), any $R \in \mathfrak{R}$ such that $\bigcap_{i \in N} R_i \subseteq R$ includes $\mathcal{R}$ as a subset. So L further constrains $Ext(\mathcal{R})$. □

So, L is an intermediate position between $\{R_i : i \in N\}$ and $Ext(\mathcal{R})$ (or $Ext(\bigcap_{i \in N} P_i)$). By preserving categorical weak preference, elements of L generally disagree with the Pareto-extension rule which requires converting any categorical weak preference into indifference unless the preference is also categorically strict. Moreover, since $\bigcap_{i \in N} I_i \subseteq \bigcap_{i \in N} R_i$, we have immediately that $\bigcap_{i \in N} I_i \subseteq \bigcap L$. That is, judgments of categorical indifference are preserved by L. Finally, with very minimal revision to the proofs below, the corresponding possibility results (cf. Propositions 25 and 26) can be stated and proved.

Proof of Propositions 25 and 26

We show something stronger than either Proposition 25 or Proposition 26 from which both results follow: $\mathcal{F}(R_1, ..., R_n) = Ext(\mathcal{R})$ satisfies *all* of the conditions involved in both Arrow's and Sen's results (and the variant generalization of **IIA** that swaps the existential quantifiers for universal quantifiers to boot).

Proof. Let $\mathcal{F}(R_1, ..., R_n) = Ext(\mathcal{R})$. We check each axiom under the assumptions that $\mathcal{F}$ satisfies **U** and $|X| > 2$.

(**P**). Suppose xP_iy for all $i \in N$, i.e., $(x, y) \in \bigcap_{i \in N} P_i$. So, $(x, y) \in \mathcal{R} \backslash \mathcal{R}^{-1}$. Since $\mathcal{R} \backslash \mathcal{R}^{-1} \subseteq R \backslash R^{-1}$ for all $R \in Ext(\mathcal{R})$, we have $(x, y) \in R \backslash R^{-1}$ for all $R \in \mathcal{F}(R_1, ..., R_n)$.

(**IIA**). Let $(R_1, ..., R_n), (R'_1, ..., R'_n) \in D$, and assume that xR_iy *iff* xR'_iy for all $i \in N$. We need to show that xRy for some $R \in Ext(\mathcal{R})$ *iff* $xR'y$ for some $R' \in Ext(\mathcal{R}')$. So, for the first direction, let xRy for some $R \in Ext(\mathcal{R})$. It follows that $(y, x) \notin \mathcal{R} \backslash \mathcal{R}^{-1}$, i.e., $(y, x) \notin \bigcap_{i \in N} P_i$. So xR_iy for some $i \in N$. It follows from the assumption that xR'_iy for the same $i \in N$. By Proposition 24, there is some $R' \in Ext(\mathcal{R}')$ such that $xR'y$ since $R'_i \in Ext(\mathcal{R}')$. An analogous argument establishes the other direction.

(**IIA′**). Let $(R_1, ..., R_n), (R'_1, ..., R'_n) \in D$, and assume that xR_iy *iff* xR'_iy for all $i \in N$. We need to show that xRy for *all* $R \in Ext(\mathcal{R})$ *iff* $xR'y$ for *all* $R' \in Ext(\mathcal{R}')$. For the first direction, suppose that xRy for all $R \in Ext(\mathcal{R})$. Then, $(x,y) \in \mathcal{R}$ (otherwise, there would be some extension of $\mathcal{R}$ that did not include (x,y)). So, either $(x,y) \in \bigcap_{i\in N} P_i$ or $(x,y) \in \bigcap_{i\in N} I_i$. Suppose that $(x,y) \in \bigcap_{i\in N} P_i$. It follows from our assumption that $(x,y) \in \bigcap_{i\in N} P'_i$. Hence, $(x,y) \in \mathcal{R}'$. So, $(x,y) \in R'$ for every $R' \in Ext(\mathcal{R}')$. Now suppose that $(x,y) \in \bigcap_{i\in N} I_i$. Then, $(x,y) \in R_i$ for all $i \in N$. By the assumption, it follows that $(x,y) \in R'_i$ for all $i \in N$. But (y,x) must also be in $\bigcap_{i\in N} I_i$. Hence, $(y,x) \in R'_i$ for all $i \in N$, too. So $(x,y) \in \bigcap_{i\in N} I'_i$. It follows that $(x,y) \in \mathcal{R}'$. So, $(x,y) \in R'$ for all $R' \in Ext(\mathcal{R}')$.

(**D**). By Proposition 24, $\{R_i : i \in N\} \subseteq Ext(\mathcal{R})$. So there is no $i \in N$ such that, for every $x,y \in X$ and every $(R_1, ..., R_n) \in D$ xP_iy implies xPy for all $R \in \mathcal{F}(R_1, ..., R_n)$ since $\neg xPy$ for some $R \in \mathcal{F}(R_1, ..., R_n)$ for a profile $(R_1, ..., R_n) \in D$ such that $\neg xP_jy$ for some $j \in N$.

(**L**). By Proposition 24, $\{R_i : i \in N\} \subseteq Ext(\mathcal{R})$. So, for each $i \in N$, if xP_iy, then there is some $R \in Ext(\mathcal{R})$ such that $(x,y) \in R \backslash R^{-1}$. □

Choice Functions

As observers, we might worry that we do not have access to an agent's "internal" preferences. Her choices, however, are more easily observed. The guiding idea in revealed preference theory is that an agent's choices reveal her preferences. There are a number of ways to formalize this idea. The *base relation* R_{C_b} is given by setting $xR_{C_b}y$ iff $x \in C(\{x,y\})$. The *Samuelson relation* R_{C_s} is given by setting $xR_{C_s}y$ iff $x \in C(Y)$ for some $Y \in \mathcal{X}$ such that $y \in Y$.

The following are some of the central properties in the study of choice functions, sometimes called *coherence* or *consistency* constraints. Many different choice-determining structures can be characterized by different combinations of them.

(α) If $S \subseteq T$, then $S \cap C(T) \subseteq C(S)$

(γ) $\bigcap_{j\in J} C(S_j) \subseteq C(\bigcup_{j\in J} S_j)$

(β) If $S \subseteq T$, $x, y \in C(S)$, and $x \in C(T)$, then $y \in C(T)$

(Aiz) If $S \subseteq T$ and $C(T) \subseteq S$, then $C(S) \subseteq C(T)$

Below, we assume that $\mathcal{X} = \mathscr{P}_{fin}(X) \backslash \{\varnothing\}$, the set of finite, non-empty subsets of X.

Proposition 33. *Let $C : \mathcal{X} \to \mathcal{X}$ be a choice function on $\mathcal{X}$. If C satisfies Property α, then $R_{C_b} = R_{C_s}$.*

Proof. Suppose that C satisfies α. Let $(x, y) \in R_{C_b}$, i.e., $x \in C(\{x, y\})$. It follows immediately that there is a $Y \in \mathcal{X}$ such that $y \in Y$ and $x \in C(Y)$. So, $(x, y) \in R_{C_s}$. Now let $(x, y) \in R_{C_s}$, i.e., there is some $Y \in \mathcal{X}$ such that $y \in Y$ and $x \in C(Y)$. By α, it follows that $x \in C(\{x, y\})$. So, $(x, y) \in R_{C_b}$. □

A binary relation, R, is called *acyclic* if, for every $n \in \mathbb{Z}^+$ and all $x_1, x_2, ..., x_n \in X$, if $x_1 P x_2, x_2 P x_3, ..., x_{n-1} P x_n$, then $\neg x_n P x_1$. So preference relations are acyclic if there are no cycles in *strict* preference. Samuelson revealed preference is by definition acyclic.

Proposition 34. *For any choice function, $C : \mathcal{X} \to \mathcal{X}$, R_{C_S} is acyclic.*

Proof. Suppose that $x_1 P_{C_s} x_2, ..., x_{n-1} P_{C_s} x_n$. Consider the menu $Y = \{x_j : j = 1, ..., n\}$. Since $x_j P_{C_s} x_{j+1}$ for $j = 1, ..., n-1$, we have $\neg x_{j+1} R_{C_s} x_j$. Since $C(Y) \neq \varnothing$, it follows that $C(Y) = \{x_1\}$. Hence, $\neg x_n P_{C_s} x_1$. □

From Propositions 33 and 34, the next proposition immediately follows.

Proposition 35. *If C is a choice function on $\mathcal{X}$ satisfying Property α, then R_{C_b} is acyclic.*

Proposition 36. *Let C be a choice function on $\mathcal{X}$. C is rationalizable iff it is rationalizable by R_{C_b}.*

Proof. (⇒) Suppose that C is rationalizable, i.e., there is an R on X such that, for all $Y \in \mathcal{X}$, $C(Y) = \{x \in Y : xRy \text{ for all } y \in Y\}$. We have to show that $R = R_{C_b}$. Suppose xRy. For *reductio*, suppose $\neg xR_{C_b}y$. So $x \notin C(\{x, y\})$. But setting $Y = \{x, y\}$, since xRy, $x \in C(\{x, y\})$, which is a contradiction. Suppose that $xR_{C_b}y$ and $\neg xRy$. So, $x \in C(\{x, y\})$. Again, setting $Y = \{x, y\}$, it follows from the assumption that $x \notin C(\{x, y\})$, which is a contradiction. Hence, $R = R_{C_b}$.

(⇐) Trivial. □

Theorem 21. *(e.g., Sen, 1971, Theorem 9) Let C be a choice function on $\mathcal{X}$. C is rationalizable iff it satisfies α and γ.*

Theorem 22. *(e.g., Sen, 1977, Proposition 11) Let C be a choice function on $\mathcal{X}$, the set of all finite subsets of X. C is weak order rationalizable iff it satisfies α and β.*

If no appeal is made to second (third, etc.) tier preferences, then any choice function generated by any set of orderings on X by V-admissibility satisfies property α. Moreover, it satisfies Aizerman's axiom, *Aiz*.

Proposition 37. *Let $\mathcal{O}$ be a set of weak orderings of X. Let $C : \mathcal{X} \to \mathcal{X}$ be defined by $C_{\mathcal{O}}(Y) = \{x \in Y : \exists R \in \mathcal{O}\ \forall y \in Y\ xRy\}$. Then, C satisfies α and Aiz.*

Proof. Suppose that $S \subseteq T$ and let $x \in S \cap C_{\mathcal{O}}(T)$. So, there is some $R \in \mathcal{O}$ such that xRy for all $y \in Y$. But since $S \subseteq T$, xRy for all $y \in C_{\mathcal{O}}(S)$ also. For *Aiz*, suppose $S \subseteq T$ and $C_{\mathcal{O}}(T) \subseteq S$. Let $x \in C_{\mathcal{O}}(S)$. Then, there is an $R \in \mathcal{O}$ such that xRy for all $y \in S$. In particular, xRy for all $y \in C_{\mathcal{O}}(T) \subseteq S$. Let t^* be an R-optimal option in T, i.e., t^*Rt for all $t \in T$. Hence, $t^* \in C_{\mathcal{O}}(T)$. So, xRt^*; therefore, $x \in C_{\mathcal{O}}(T)$. □

Recall that generating a choice function from a set of orderings by V-admissibility does not secure the rationalizability of the resulting choice function.

Example 4. *Let* $X = \{x, y, z\}$ *and consider* $\mathcal{O} = \{R_1, R_2\}$ *with* xP_1yP_1z *and* zP_2yP_2x. *(Perhaps the professor changed some goals in her private life.) Using* V*-admissibility, for non-singleton sets we have* $C(\{x, y\}) = \{x, y\}$, $C(\{y, z\}) = \{y, z\}$, $C(\{x, z\}) = \{x, z\}$, *and* $C(\{x, y, z\}) = \{x, z\}$. *Since* $y \in C(\{x, y\}) \cap C(\{y, z\})$, *but* $y \notin C(\{x, y, z\})$, C *violates* γ. *It follows from Theorem 21 that* C *is not rationalizable by a binary relation.*

Choice structures like V-admissibility have been studied in the literature. For finite X and domain of choice consisting of all non-empty subsets of X, Aizerman and Malishevski characterized *pseudo-rationalizability* also called *collected extremal choice*. C is pseudo-rationalizable if there exists a finite set of orderings $\{R_i : i = 1, ..., n\}$ such that, for all $Y \in \mathcal{X}$, the choice function selects the R_i-maximal elements for $i = 1, ..., n$:

$$C(Y) = \bigcup_{1 \leqslant i \leqslant n} max\ R_i \cap Y \times Y$$

Theorem 23. *(Aizerman and Malishevski, 1981, Theorem 3; Moulin, 1985, Theorem 5) Let* X *be finite and let* C *be a choice function on* $\mathcal{X}$. C *is pseudo-rationalizable iff it satisfies* α *and Aiz.*

Pedersen generalizes this result, relaxing both the assumption that X is finite and that the set of menus is restricted to $\mathcal{X}$ (2009a).

Let M be the set of *undominated* options in X according to a set of weak orderings, $\mathcal{O}$:

$$M = \{x \in X : \forall y \in X\ \exists R \in \mathcal{O}\ xRy\}$$

Instead of looking at only the undominated options in all of X, we can consider the undominated options in a particular subset of X, which may or may not be dominated in X.

$$M_Y = \{x \in Y : \forall y \in Y\ \exists R \in \mathcal{O}\ xRy\}$$

Of particular interest is the set of undominated options according to $\mathcal{O} = Ext(\mathcal{R})$, where $\mathcal{R}$ is categorical preference for R_i, $i \in N$.

Proposition 38. *Let M_Y be the undominated options in Y according to $Ext(\mathcal{R})$. There is a weak order extension of $\mathcal{R}$ (i.e., an $R \in Ext(\mathcal{R})$) such that, for all $x \in M_Y$, xRy for all $y \in Y$.*

Proof. Set $\mathcal{R}_{M_Y} = \mathcal{R} \cup (M_Y \times Y)$. We show first that $\mathcal{R}_{M_Y}$ is a quasiorder. That $\mathcal{R}_{M_Y}$ is reflexive follows from the fact that $\Delta \subseteq \mathcal{R} \subseteq \mathcal{R}_{M_Y}$. Suppose that $x\mathcal{R}_{M_Y}y$ and $y\mathcal{R}_{M_Y}z$. There are four cases to consider to establish $xR_{M_Y}z$.

(1) $(x,y),(y,z) \in \mathcal{R}$. By Proposition 22, $\mathcal{R}$ is transitive, so $(x,z) \in \mathcal{R}$. Hence, $(x,z) \in \mathcal{R}_{M_Y}$.

(2) $(x,y),(y,z) \in M_Y \times Y$. Then, $x \in M_Y$. It follows from the construction of $\mathcal{R}_{M_Y}$ that $(x,z) \in \mathcal{R}_{M_Y}$.

(3) $(x,y) \in \mathcal{R}$ and $(y,z) \in M_Y \times Y$. Either $(x,y) \in \bigcap_{i \in N} P_i$ or $(x,y) \in \bigcap_{i \in N} I_i$. Suppose the former. Then, yRx for no $R \in Ext(\mathcal{R})$. So $y \notin M_Y$, which contradicts our assumption that $(y,z) \in M_Y \times Y$. Suppose instead that $(x,y) \in \bigcap_{i \in N} I_i$. Also suppose, for reductio, that $x \notin M_Y$. Then y is not in M_Y either (since any option that dominates x dominates y), which contradicts our assumption. So, $x \in M_Y$ and $x\mathcal{R}_{M_Y}z$.

(4) $(y,z) \in \mathcal{R}$ and $(x,y) \in M_Y \times Y$. Then, $x \in M_Y$, so $x\mathcal{R}_{M_Y}z$.

Now we need to show that $\mathcal{R}_{M_Y}$ is a quasiorder extension of $\mathcal{R}$. Clearly $\mathcal{R} \subseteq \mathcal{R}_{M_Y}$. Let $(x,y) \in \mathcal{R}\backslash\mathcal{R}^{-1} = \bigcap_{i \in N} P_i$. Suppose for reductio that $(y,x) \in \mathcal{R}_{M_Y}$. Then it must be the case that $y \in M_Y$. But then yRx for some $R \in Ext(\mathcal{R})$. This contradicts the fact that $\mathcal{R}\backslash\mathcal{R}^{-1} \subseteq R\backslash R^{-1}$ for all $R \in Ext(\mathcal{R})$. So $(y,x) \notin \mathcal{R}_{M_Y}$. Hence, $\mathcal{R}_{M_Y}$ is a quasiorder extension of $\mathcal{R}$. By Theorem 19, $\mathcal{R}_{M_Y}$ has a weak order extension, R, which in turn is a weak order extension of $\mathcal{R}$. Furthermore, since $M_Y \times Y \subseteq R$, for all $x \in M_Y$ and all $y \in Y$ we have xRy. □

Proposition 39. *(Cf. Levi, 1986a, pp. 100-101) $C_{Ext(\mathcal{R})}$ satisfies γ.*

Proof. Let $x \in \bigcap_{j \in J} C_{Ext(\mathcal{R})}(S_j)$. So, for every $y \in \bigcup_{j \in J} S_j$, there is an $R \in Ext(\mathcal{R})$ such that xRy. So x is undominated in $\bigcup_{j \in J} S_j$: $x \in M_{\bigcup_{j \in J} S_j}$. By Proposition 38, there is an $R \in Ext(\mathcal{R})$ such that xRy for all $y \in \bigcup_{j \in J} S_j$. So, $x \in C_{Ext(\mathcal{R})}(\bigcup_{j \in J} S_j)$. □

Together, Propositions 37 and 39 and Theorem 21 imply that there exists a single binary relation that rationalizes $C_{Ext(\mathcal{R})}$. In this sense, requiring that all compatible extensions of categorical preference be permissible restores binariness. In fact, we can be more specific.

Proposition 40. *Let R_{Pe} be the Pareto-extension relation for $(R_1, ..., R_n)$ and R_{C_b} be base revealed preference for $C_{Ext(\mathcal{R})}$. $R_{Pe} = R_{C_b}$.*

Proof. By Proposition 36 and the fact that $C_{Ext(\mathcal{R})}$ is rationalizable, we have that $C_{Ext(\mathcal{R})}$ is rationalized by R_{C_b}. First, assume that $xR_{C_b}y$. Then, $(y, x) \notin \bigcap_{i \in N} P_i$. Hence, $xR_{Pe}y$. Next, assume that $xR_{Pe}y$. Either $(x, y) \in \mathcal{R}$ or not. Suppose the former. Then, there is some $R \in Ext(\mathcal{R})$ such that xRy. So, $x \in C_{Ext(\mathcal{R})}(\{x, y\})$. Hence, $xR_{C_b}y$. Suppose on the other hand that $(x, y) \notin \mathcal{R}$. Since $xR_{Pe}y$, $(y, x) \notin \bigcap_{i \in N} P_i$. So, there is a quasiorder extension of $\mathcal{R}$ that includes (x, y), which in turn has a weak order extension. So, again, $x \in C_{Ext(\mathcal{R})}(\{x, y\})$; therefore, $xR_{C_b}y$. □

Propositions 25 may not be as surprising given Proposition 40 since it is known that, in general, relaxing ordering allows certain possibilities to be established and, in particular, the Pareto-extension rule satisfies Arrow's conditions (e.g., Sen, 1977, Proposition 1). But the generalized framework studied in this essay can still mark distinctions between the Pareto-extension rule and the $Ext(\mathcal{R})$ model. Sen established that neither the Pareto-extension rule nor any other binary relation generating a choice function can satisfy **U**, **P**, and **L**. For our generalizations of those axioms, $Ext(\mathcal{R})$ *does* satisfy them, while the Pareto-extension rule still does *not*. Moreover, we have given a number of reasons for general philosophical reservations about revealed preference. In the first place, (base) revealed preference does not generally coincide with judgments of admissibility in the presence of indeterminacy (e.g., Example 4). In the second place are the conceptual and widely-appreciated concerns about revealed preference theory rehearsed in Section 9.2. For example, from the point of view of indeterminacy, the Pareto-extension rule often illicitly converts incomparability into indifference.