Ergo AN OPEN ACCESS JOURNAL OF PHILOSOPHY

Regret Averse Opinion Aggregation

LEE ELKIN
Jagiellonian University in Kraków

It is often suggested that when opinions differ among individuals in a group, the opinions should be aggregated to form a compromise. This paper compares two approaches to aggregating opinions, linear pooling and what I call opinion agglomeration. In evaluating both strategies, I propose a pragmatic criterion, No Regrets, entailing that an aggregation strategy should prevent groups from buying and selling bets on events at prices regretted by their members. I show that only opinion agglomeration is able to satisfy the demand. I then proceed to give normative and empirical arguments in support of the pragmatic criterion for opinion aggregation, and that ultimately favor opinion agglomeration.

The chairman of an agricultural committee has commissioned a group of economic experts to determine how likely the Food Price Index (FPI) will exceed 120 by 2025, given the rising number of droughts annually (the monthly FPI average for 2020 is near 95). After reviewing the data, the members exchange their opinions on the matter, but to no one's surprise, the opinions differ. The group realizes, however, that in moving forward, they need to come to a reasonable compromise, if not a consensus, before briefing the chairman, but how?[1] It is often suggested that such a collective problem can be solved by aggregating the opinions. In this paper, I will compare two strategies for aggregating opinions, the popular linear pooling approach and what I call opinion agglomeration.

The way this paper differs from previous work on the subject is by its focus on a pragmatic dimension of opinions, following de Finetti (1974), where a

1. The difference between a consensus and a compromise is that the former requires individuals to continuously deliberate and revise their opinions until they reach an agreement, whereas the latter requires individuals to find common ground for collective decision-making purposes, but without necessarily changing their opinions (see Wagner 2009; Moss 2011).

Contact: Lee Elkin <lee.elkin@uj.edu.pl>

https://doi.org/10.3998/ergo.1153

probabilistic opinion of an event *X* is viewed as the maximum price one is willing to buy and minimum price they are willing to sell a bet that pays 1 monetary unit in case *X* occurs, nothing otherwise. After extending this pragmatic interpretation to collective opinions, I will propose a novel constraint, entailing that an aggregation strategy should prevent groups from buying and selling bets on events at prices *regretted* by their members. The felt regret in question is disappointment over paying too much for a bet and similarly selling the bet for too little. The regret is thus *ex ante* rather than *ex post*, as the negative emotion is conditioned on acts that precede learning the outcomes that settle the bets.

With the pragmatic constraint imposed on opinion aggregation, I will show that linear pooling fails to satisfy the demand when opinions differ provided that a group's losses from buying and selling bets at the prices yielded by the aggregation strategy are perceived to be positive by some members, thus causing those members to regret the group's betting behavior. At best, linear pooling may eliminate the perceived loss from one of the two transaction types for all members, but only if the Non-Dictatorship criterion is violated and the perceived loss from the opposite transaction type is increased from the viewpoints of some members. In comparison, I will show that opinion agglomeration satisfies the demand, as a group's losses from buying and selling bets at the prices yielded by the aggregation strategy are perceived to be zero by all members. Thus, no member regrets the group's betting behavior. Afterward, I will provide normative and empirical arguments in support of the pragmatic criterion for opinion aggregation, and that ultimately favor opinion agglomeration.

1. Two Aggregation Strategies

While there are many approaches to aggregating opinions (see Genest & Zidek 1986; Dietrich & List 2016), I will restrict the focus to just two in this paper, linear pooling and opinion agglomeration. In this section, I will introduce the technical details, along with reasons for why groups might adopt each strategy.

1.1. Linear Pooling

For a group of individuals $i = 1,\dots,n$, let the set of events the group is concerned with be a finite algebra $\mathcal{A}$ over a set of states Ω. Call $\mathcal{A}$ the group's *agenda*. With respect to the agenda, every individual i's opinions are represented by a function, $p_i : \mathcal{A} \to [0,1]$, and $p_i \in \Delta$, where Δ is the set of all probability functions on $\mathcal{A}$. Call a collection of opinion functions for all i, $(p_1,\dots,p_n)$, an *opinion profile*, and let Δ^n be the set of admissible opinion profiles. Finally, let the group's opinions be formed by a *pooling function*, F, mapping opinion profiles to Δ.

The most common linear pooling model, and the one I will focus on in this paper, is the following:

$$F_L(p_1,\ldots,p_n)(X)=\sum_{i=1}^{n} w_i p_i(X) \quad \forall X \in \mathcal{A}, \tag{1}$$

where w_i is a non-negative weight representing individual i's level of reliability or expertise, and the sum of weights for all i is equal to one (Stone 1961). Simply put, the linear pooling function, F_L, yields weighted averages of individual opinions for all events X that represent the opinions of the group.

The pooling model introduced may already be familiar to those acquainted with the epistemic peer disagreement literature, as it is a generalization of the so-called *equal weight* model (see Jehle & Fitelson 2009). On an equal weight view, peers should resolve their disagreements through the following:

$$F_{ew}(p_1,\ldots,p_n)(X)=\frac{1}{n}\sum_{i=1}^{n} p_i(X) \quad \forall X \in \mathcal{A}. \tag{2}$$

Besides settling peer disagreements, (2) is an efficient way of forming collective opinions without having to invest time in determining the weight given to each individual. But in case the group is committed to determining how much weight each individual deserves, and every individual is given a say on the matter, the group may face yet another problem, namely, disagreement over weighting assignments. To address the problem, some suggest that the individuals should revise their opinions to weighted averages and iterate the procedure until all opinions have sufficiently converged (DeGroot 1974; Lehrer 1976; Wagner 1978).

The formal setup is a bit different in the DeGroot-Lehrer-Wagner tradition. Let w_{ij}, for $i, j = 1,\ldots,n$, be the weight that individual *i* assigns to every individual *j*, including *i*, and the sum of weights is equal to one. These weights reflect personal opinions about every individual's level of reliability or expertise from *i*'s point of view and are situated as a row in a Markov matrix,

$$\mathbf{M}=\begin{bmatrix} w_{11}, & \cdots & ,w_{1n} \\ . & \cdots & . \\ . & \cdots & . \\ . & \cdots & . \\ w_{n1}, & \cdots & ,w_{nn} \end{bmatrix}, \tag{3}$$

with the opinion profile situated as a column vector,

$$\mathbf{P} = \begin{bmatrix} p_1 \\ \cdot \\ \cdot \\ \cdot \\ p_n \end{bmatrix}. \tag{4}$$

The opinions are then aggregated repeatedly through iterated multiplication, and as DeGroot (1974) showed, $\mathbf{P}^k$ defined as $\mathbf{M}^k\mathbf{P}$ converges to some opinion profile $\mathbf{P}^\infty$ such that no subsequent revision will change the limiting profile.[2]

Put in more familiar terms, every individual i revises at step k to the following:

$$p_i^k(X) = \sum_{j=1}^{n} w_{ij} p_j^{k-1}(X) \quad \forall X \in \mathcal{A}. \tag{5}$$

As k approaches infinity, opinions will sufficiently converge, and no further revision will change any one opinion in the profile. The upshot of adopting this iterative procedure is that the group will arrive at a consensus as opposed to a mere compromise. The underlying idea is that deliberation compels the individuals to change their opinions, and by continuously repeating the process, the group will eventually come to an agreement. Unfortunately, the consensus approach with varying subjective weight distributions can be costly timewise.

If, however, the individuals show *mutual respect*, that is, every individual allocates equal weight to all members of the group, then they will settle on the following:

$$p_i^1(X) = \sum_{j=1}^{n} \frac{1}{n} p_j^o(X) \quad \forall X \in \mathcal{A}. \tag{6}$$

Notice that mutual respect is the most efficient solution since a consensus is reached after the first deliberation, thus making it computationally appealing.

What should be apparent to the reader at this point is that linear pooling consists of a class of pooling models (larger than the two mentioned above) capable of forming compromises for different occasions. Besides promoting a compromise, I illustrated how weighted averaging also leads to a consensus through iterated belief revision. Considering the benefits afforded by weighted averaging, pooling by F_L is often viewed as an attractive aggregation strategy.

2. Note that the result depends on further assumptions. See Theorem 2 in DeGroot (1974).

1.2. Motivating Linear Pooling

There are at least two reasons for favoring linear pooling. The first comes from discussions on peer disagreement (Christensen 2007; Elga 2007; Feldman 2007). The popular *conciliatory* view maintains that peers should move their opinions toward one another after discovering that they disagree. Proponents contend that a difference of opinion is evidence that each peer is liable for having made a mistake provided that no peer is assumed to be epistemically superior to another. Such evidence is said to undermine the reliability of each peer's reasoning on the matter. And since there is no good epistemic reason to dismiss a peer's opinion, peers should welcome the evidence and take the middle ground for the time being. While there is debate over what the correct belief revision amounts to (see Christensen 2011; Kelly 2013), many have suggested that peers should adopt an equal weighted average, as it seems to be the most intuitive answer.

The second relates to "the wisdom of crowds" (Surowiecki 2005). In short, the wisdom of crowds is the idea that groups outperform individuals in forecasting and estimation tasks (under certain conditions such as the independence of every opinion) and has been confirmed empirically. One of the earliest observations traces back to Sir Francis Galton (1907), who reported that the median guess of the weight of a "dressed" ox in a weight judging contest was within less than one percent of the ox's true weight of 1198 lb. Although Galton highlighted the accuracy of the median in the original paper, R. H. Hooker (1907) pointed out in a reply that the arithmetic mean of 1196 lb, based on the numbers reported, was more accurate than the median. In his response to Hooker, Galton provided the correct arithmetic mean of 1197 lb (only 1 lb off the mark) versus the 1207 lb median, thereby confirming Hooker's suspicion of the mean being more accurate than the median.[3] Despite the ensuing dispute, what was discovered in the exchange is the value of group estimates and predictions when opinions are averaged. Thus, the wisdom of crowds provides a further reason for adopting at least a special case of linear pooling given the accuracy of the average opinion.

1.3. Agglomerative Pooling

The second strategy I will consider and call opinion agglomeration requires a more expressive imprecise probability (IP) framework (Moral & del Sagrado 1998; Nau 2002; Stewart & Quintana 2018). For a group of individuals $i = 1, \ldots, n$, let the group's agenda once again be a finite algebra $\mathcal{A}$ over a set of states Ω.

3. Galton remained committed to his preference for the median, despite confirming that the arithmetic mean turned out to be a more accurate estimate.

With respect to the agenda, every individual i's opinions are represented by a non-empty, convex set of probability functions, $P_i \subseteq \Delta$, on the agenda.[4] Call a collection of set-based opinions for all i, $(P_1,\ldots,P_n)$, an opinion profile, and let $\mathcal{D}^n$ be the set of admissible opinion profiles, where $\mathcal{D}$ is the set of all non-empty, convex sets of probability functions. Finally, let the group's opinions be formed by a pooling function, $\mathcal{F}$, mapping opinion profiles to $\mathcal{D}$.

The most natural agglomerative pooling model, and the one I will focus on in this paper, is the following:

$$\mathcal{F}_A\left(P_1,\ldots,P_n\right)\left(X\right) = \text{conv}\left(\cup_i P_i\left(X\right)\right) \qquad \forall X \in \mathcal{A}. \tag{7}$$

Simply put, the agglomerative pooling function, $\mathcal{F}_A$, yields non-empty, convex sets of probabilities for all events X, representing the opinions of the group.[5] Despite the absence of a reliability or expertise factor, a form of equal weight is encoded, as the function's output is invariant under permutations of individuals.

While pooling by $\mathcal{F}_A$ may come across as a sensible way of agglomerating opinions, some might question the convexity property provided that there are cases in which non-convex sets of probability functions are intuitive. For example, if a group were to learn that a coin either has an exact bias of 0.6 or 0.7 toward heads, then it would be wise to adopt a set consisting of probability functions assigning just those two values to H rather than a set that also includes the convex hull, as the latter encodes additional information. Thus, it seems that convexity should be optional for agglomerative pooling, not mandatory.

Let me clarify that I only assume the convexity property here for mathematical convenience and make no philosophical commitment to it like Levi (1974) and Joyce (2010). But in case one insists on relaxing the convexity assumption, they might prefer the following agglomerative pooling model instead:

$$\mathcal{F}_A'\left(\mathbb{P}_1,\ldots,\mathbb{P}_n\right)(X) = \cup_i \mathbb{P}_i(X) \qquad \forall X \in \mathcal{A}, \tag{8}$$

4. Note that 'imprecise probability' is conventionally used as an umbrella term that not only covers sets of probability functions, but also belief functions (Shafer 1976), Choquet capacities (Wasserman & Kadane 1990), possibility measures (Dubois & Prade 1988), lower previsions (Walley 1991), and sets of desirable gambles (Walley 1991; Couso & Moral 2009). Although sets of probability functions are not the most general, they subsume belief functions and possibility measures and thus are not the least general. The reason for choosing sets of probability functions here, however, is merely to follow the recent trend in epistemology (see Bradley 2019 for a survey on IP in philosophy).

5. Stewart and Quintana (2018) suggest a rule along these lines as a way of generalizing the IP pooling rule they focus on in their paper: conv $\{p_i : i = 1,\ldots,n\}$.

where $\mathbb{P}$ denotes a set of probability functions that is not necessarily convex (note that the pooling function, $\mathcal{F}_A^{'}$, has a domain $\mathcal{P}(\Delta)^n$ and range $\mathcal{P}(\Delta)$, where the latter denotes the power set of delta). One advantage of the $\mathcal{F}_A^{'}$ agglomerative pooling function is that it preserves judgments about the independence of events. That is, if all members of a group judge events X and Y to be stochastically independent, $\mathcal{F}_A^{'}$ preserves their independence in the collective opinions, whereas $\mathcal{F}_A$ does not. There has been much debate over this issue in the literature (see, e.g., Levi 1974; Kyburg & Pittarelli 1996), which goes beyond the scope of this paper, but I have introduced (8) for those concerned about convexity.

As it should be apparent by now, opinion agglomeration, like linear pooling, consists of a class of pooling models (larger than the two mentioned above) capable of forming compromises for different occasions. Although the agglomerating method may not have the same level of appeal as weighted averaging, the pooling function, $\mathcal{F}_A$, enjoys some nice properties, as we will come to see.

1.4. Motivating Opinion Agglomeration

A reason for modeling opinions by imprecise probabilities is that individuals occasionally form non-additive beliefs, as notably demonstrated by Daniel Ellsberg (1961). In surveying decision theorists, one of the problems given involved bets on drawing balls from two distinct urns. Participants were informed that urn 1 contains a mixture of 50 red balls and 50 black balls, and urn 2 contains a mixture of 100 red and black balls in unknown proportions. They were asked to report their preferences for a series of bets with equal prizes on randomly drawing a red ball and drawing a black ball. Between two bets on a single urn, most were indifferent. Between two bets across urns, most reported a strict preference for the bet on urn 1, e.g., a $100 or nothing bet on drawing a red ball from urn 1 was strictly preferred to a $100 or nothing bet on drawing a red ball from urn 2.

The phenomenon described is known as *ambiguity aversion*. As Ellsberg pointed out, a strict preference for a bet on urn 1 to a bet on urn 2 cannot be characterized by the traditional Savage axioms when indifferent toward bets on each urn individually. For such preferences on L. J. Savage's (1954) view are inconsistent and thus irrational. But the followers of Herbert Simon disagree. Simon (1955) claimed that the optimality of judgments and decisions should be relativized to an individual's cognitive capacities and available information. Under this view, judgment and decision optimality should be thought of instead as *bounded rationality* (see Wheeler 2020 for a comprehensive survey). In following this tradition, Gilboa and Schmeidler (1989) introduced a set of axioms, entailing a maximin expected utility decision rule, whereby a decision maker chooses an option that maximizes minimum expected utility, relative to a convex set of probability functions. Provided the extended decision framework developed by Gilboa and Schmeidler, the reported preferences resulting from the Ellsberg experiments can plausibly be explained (see Machina & Siniscalchi 2014 and Gilboa & Marinacci 2016 for recent surveys on ambiguity aversion).

Although ambiguity aversion might convey a need for imprecise probabilities, the observed behavior only offers a reason to think that IP is useful for modeling the judgments and decisions of individuals.[6] However, Keller, Sarin, and Sounderpandian (2007) discovered that groups exhibit ambiguity aversion also. As part of their experiment, they tested participants on an Ellsberg-like problem, first by one's self and then collectively in randomly formed dyads. In addition to replicating the findings of Ellsberg, they found that the average amount dyads are willing to pay for a bet on the ambiguous urn is less than the average amount individuals are willing to pay for the same bet, implying that groups are even more ambiguity averse than their members. What is more, the data suggest that the opinions of dyads are some conglomerate of individual opinions given a larger difference between the buying and selling prices of groups than that of individuals, thus furnishing empirical evidence for compromising by agglomerating.

2. Desirable Criteria for Opinion Aggregation

In the previous section, I presented two strategies for aggregating opinions. In this section, I aim to assess their plausibility by comparing them against a set of desirable criteria that aggregation strategies arguably should meet. A handful of criteria for pooling functions have been proposed in the past, which I will briefly review. Then, I will introduce a pragmatic criterion called No Regrets. Distinct from the other criteria, the latter focuses on collective behavior. A reason for introducing the criterion is that, as far as I know, there has been very little to no attention given to pragmatic considerations for opinion aggregation. This deficiency is quite surprising since a difference of opinion should, among other things, prompt *cautious* group behavior if groups were to act on their opinions, which the proposed criterion aims to accommodate in some form.

6. Of course, ambiguity aversion is not the only reason for adopting imprecise probabilities. Other considerations include accommodating suspension of judgment in probabilistic models (see, e.g., Levi 1983; Sturgeon 2010), respecting evidence (see, e.g., Joyce, 2005, 2010; Schoenfield 2012; Konek in press), and vague credence (see, e.g., Lyon 2017). For comprehensive surveys on various applications of imprecise probability, see Augustin, Coolen, De Cooman, and Troffaes (2014) and Bradley (2019).

2.1. A Brief Review of Some Desirable Criteria for Pooling Functions

Although a number of criteria have been proposed in the past (see, e.g., Aczél & Wagner 1980; McConway 1981; Genest & Zidek 1986; Garg, Jayram, Vaithyanathan, & Zhu 2004; Dietrich & List 2016; Stewart & Quintana 2018; Hartmann 2017), the following four regularly appear in the literature due to the epistemic features they preserve.

(**i**) Unanimity. For all admissible opinion profiles $(p_1,\ldots,p_n)$, if all p_i are identical, then $F(p_1,\ldots,p_n)=p_i$ (in IP: for all admissible opinion profiles $(P_1,\ldots,P_n)$, if all P_i are identical, then $\mathcal{F}(P_1,\ldots,P_n)=P_i$).

(**ii**) Eventwise Independence or Weak Setwise Function Property. There exists a function $G:\mathcal{A}\times[0,1]^n\to[0,1]$ such that for all admissible opinion profiles $(p_1,\ldots,p_n)$ and events $X\in\mathcal{A}$, $F(p_1,\ldots,p_n)(X)=G(X,p_1(X),\ldots,p_n(X))$ (in IP: there exists a function $\mathcal{G}:\mathcal{A}\times\mathcal{P}([0,1])^n\to\mathcal{P}([0,1])$ such that for all admissible opinion profiles $(P_1,\ldots,P_n)$ and events $X\in\mathcal{A}$, $\mathcal{F}(P_1,\ldots,P_n)(X)=\mathcal{G}(X,P_1(X),\ldots,P_n(X))$.

(**iii**) Boundedness. For all admissible opinion profiles $(p_1,\ldots,p_n)$ and events $X\in\mathcal{A}$, $F(p_1,\ldots,p_n)(X)$ is in the closed interval [min $(p_1(X),\ldots,p_n(X))$, $\max(p_1(X),\ldots,p_n(X))$] (in IP: for all admissible opinion profiles $(P_1,\ldots,P_n)$ and events $X\in\mathcal{A}$, $\mathcal{F}(P_1,\ldots,P_n)(X)$ is a subset of the closed interval $[\inf\cup_i P_i(X),\sup\cup_i P_i(X)]$).

(**iv**) Non-Dictatorship. There is no individual i such that for all admissible opinion profiles $(p_1,\ldots,p_n)$, $F(p_1,\ldots,p_n)=p_i$ (in IP: there is no $\mathbf{i}$ such that for all admissible opinion profiles $(P_1,\ldots,P_n)$, $\mathcal{F}(P_1,\ldots,P_n)=P_i$).

The first two criteria trace back to early systematic studies on opinion pooling (Aczél & Wagner 1980; McConway 1981). In short, the Unanimity criterion implies that if all individuals hold the same opinions, then the collective opinions should not differ. The Eventwise Independence criterion implies that the collective opinion for any event $X\in\mathcal{A}$ should depend only on the individual opinions of X. Both of these criteria are straightforward and epistemically plausible.

The Boundedness criterion (Garg et al. 2004) implies that a pooled opinion should not fall outside of the range of opinions expressed by the members of the group. This criterion is reasonable, as the group agrees by default that the epistemically warranted opinion, based on the information available, is within the range of expressed opinions. Any pooling function yielding a collective opinion outside of the range respects none of the individual opinions. Furthermore, a pooling function violating Boundedness is prone to underestimating or overestimating the probabilities of events with uncertain chances since, as empirical evidence concerning collective intelligence suggests, groups "bracket the truth."[7]

The Non-Dictatorship criterion implies that no individual should determine the group's opinions no matter what. This is a reasonable constraint to impose on pooling functions, especially for precise probabilities, provided that (non-extreme) weighted averages will often be more accurate than any individual's opinions. The criterion is also plausible from an evidentialist point of view since it, like Boundedness, maintains respect for the evidence (Feldman 2005). For these reasons, the Non-Dictatorship criterion should appear to be feasible.

2.2. 'Regret Averse' Pooling

The above criteria for pooling functions aim at preserving features that serve the epistemic interests of groups, but groups should also care about preserving features that serve their practical interests. After all, why aggregate opinions if not for guiding collective behavior? Let us turn our attention then to pragmatic criteria. While many might exist, I will propose one relating to group betting behavior. The idea is that an aggregation strategy should prevent groups from buying and selling bets on events at prices *regretted* by their members.[8] In this section, I will flesh out the idea in full, but first, some preliminaries are in order.

The betting convention established by de Finetti (1974) assumes that every individual i is willing to buy and sell bets on events $X \in \mathcal{A}$ that pay 1 monetary unit if X occurs, nothing otherwise. For all individuals i and events X, let the maximum price i is willing to pay for a bet on X be represented by $p_i^-(X)$, and the minimum price i is willing to sell the bet be represented by $p_i^+(X)$.

7. Easwaran, Fenton-Glynn, Hitchcock, and Velasco (2016) defend a family of updating rules with a property they call *synergy*, resulting in opinions outside of the ranges of peer opinions. However, they make it clear that they are only concerned with individuals updating their own opinions, not opinion aggregation, but Russell, Hawthorne, and Buchak (2015) and Dietrich (2019) take up the group-level problem and defend an alternative pooling strategy, namely, geometric averaging that also has the synergistic property. In addition to synergy, the geometric approach is *externally Bayesian*, meaning that pooling and conditionalization commute. I bring these points to the reader's attention to show that Boundedness is not entirely uncontroversial, as it precludes geometric and multiplicative pooling, both of which have some valuable properties. But in defense of Boundedness, Pettigrew (2019) shows that any aggregation strategy violating the criterion is strictly worse in terms of accuracy than an aggregation strategy satisfying it from the perspective of each individual. Thus, aggregation strategies that satisfy the criterion yield more epistemic value than those that violate it (though, only in case of precise probabilities).

8. The concept of regret is not new to decision theory. Savage (1951) notably proposed the *minimax regret* decision criterion, and Loomes and Sugden (1982) exploited the notion in their influential theory of regret-based preferences as an alternative to expected utility theory.

On de Finetti's view, an individual's (precise) opinion $p_i(X)$ is their fair price for the bet such that $p_i^-(X) = p_i^+(X) = p_i(X)$. By comparison, the buying and selling prices in IP are canonically represented by a *lower probability* $P_i^-(X) = \inf P_i(X)$ and an *upper probability* $P_i^+(X) = \sup P_i(X)$, respectively. These prices need not be the same, though. Hence, P_i need not be precise, that is, a singleton set.

Regardless of whether opinions are precise or imprecise on some matter, all rational individuals should agree on at least one thing. That is, no individual should buy a bet on an event above or sell the bet below the buying and selling prices fixed by their opinion of that event, for the losses resulting from paying too much for the bet and selling it for too little are perceived to be positive at prices above and below. Betting behavior that leads to such losses is *regrettable* and should be avoided (see Elkin & Wheeler 2018 for a similar argument). In making the latter idea explicit, for all prices $u \in [0,\infty)$, events $X \in \mathcal{A}$, and individuals i, let the aforementioned type of regret resulting from buying and selling a bet on X for u be represented by $r_i^-(u)$ and $r_i^+(u)$, respectively, where $r_i^-(u) = \max\{u - p_i(X), 0\}$ and $r_i^+(u) = \max\{p_i(X) - u, 0\}$ (in IP: $R_i^-(u) = \max\{u - P_i^-(X), 0\}$ and $R_i^+(u) = \max\{P_i^+(X) - u, 0\}$, for all $u \in [0,\infty)$, $X \in \mathcal{A}$, and individuals i).[9]

But how do these functions capture the kind of regret suggested? The defined regret functions can be made sensible in the following way. Assume that utility is equal to monetary payoff and all individuals i evaluate risky acts according to expected utility theory, or maximin expected utility theory in case of imprecise probabilities. Considering a bet on an event X, buying the bet for a price $u \geq 0$ yields utility $1-u$ if X occurs, $-u$ otherwise. The value of buying the bet, *ex ante*, for a price u then is $p_i(X)(1-u) + (1-p_i(X))(-u) = p_i(X) - u$, and $P_i^-(X) - u$ in case of imprecise probabilities, as it is the minimum expected utility:

$$\begin{aligned}&\min_{p \in P_i}[p(X)(1-u) + (1-p(X))(-u)]\\ &= \min_{p \in P_i}[p(X) - u] = P_i^-(X) - u.\end{aligned}$$

The *ex ante* value of not buying the bet is, of course, 0. By opting to buy the bet for a price u, the *ex ante* regret incurred is

$$\begin{aligned}&\max\{\text{value of not buying the bet} - \text{ value of buying the bet}, 0\}\\ &= \max\{0 - (p_i(X) - u), 0\} = \max\{u - p_i(X), 0\}\\ &(\text{in IP} := \max\{0 - (P_i^-(\mathrm{X}) - u), 0\} = \max\{u - P_i^-(X), 0\}).\end{aligned}$$

9. The regret functions, properly defined, take prices, events, and opinions as arguments, but I will abuse notation by writing the functions as done above for simplicity.

On the other side, selling the bet for a price $u \geq 0$ yields utility $u-1$ if X occurs, u otherwise. The value of selling the bet, *ex ante*, for a price u then is $p_i(X)(u-1)+(1-p_i(X))(u) = u - p_i(X)$, and $u - P_i^+(X)$ in case of imprecise probabilities, as it is the minimum expected utility:

$$
\begin{aligned}
&\min_{p \in P_i}[p(X)(u-1)+(1-p(X))(u)] \\
&= \min_{p \in P_i}[u - p(X)] = u + \inf(-P_i(X)) \\
&= u - \sup P_i(X) = u - P_i^+(X).
\end{aligned}
$$

The *ex ante* value of not selling the bet is, of course, 0. By opting to sell the bet for a price u, the *ex ante* regret incurred is

$$
\begin{aligned}
&\max\{\text{value of not selling the bet} - \text{value of selling the bet}, 0\} \\
&= \max\{0-(u-p_i(X)),0\} = \max\{p_i(X)-u,0\} \\
&(\text{in IP: } = \max\{0-(u-P_i^+(X)),0\}) = \max\{P_i^+(X)-u,0\}).
\end{aligned}
$$

It should be clear at this point how the defined regret functions capture the (*ex ante*) regrettable losses from buying a bet for a price above one's maximum buying price and selling the bet for a price below one's minimum selling price.[10]

Suppose now that in addition to having a preference for not making bets that are regretted, *ex ante*, by oneself, the individuals similarly have a preference for the group not making bets that are regretted, *ex ante*.[11] To satisfy the latter preference for all individuals *i*, the pooling function F (or $\mathcal{F}$) adopted should yield collective opinions such that no member regrets the group buying a bet on an event X at the group's maximum buying price $F(p_1,\ldots,p_n)^-(X) = \min F(p_1,\ldots,p_n)(X)$ and selling the bet at the group's minimum selling price $F(p_1,\ldots,p_n)^+(X) = \max F(p_1,\ldots,p_n)(X)$ (in IP: $\mathcal{F}(P_1,\ldots,P_n)^-(X) = \inf \mathcal{F}(P_1,\ldots,P_n)(X)$ and $\mathcal{F}(P_1,\ldots,P_n)^+(X) = \sup \mathcal{F}(P_1,\ldots,P_n)(X)$). Put precisely,

10. I owe many thanks to an anonymous referee for their valuable suggestions on justifying the functional forms of the regret functions.

11. Under this assumption, the individuals have a common interest in preventing the group from losing money, but that commonality does not mean that they completely identify with the group, for they might differ in opinion. Thus, it is supposed that the individuals have a common interest, but may hold conflicting opinions. In many economic models, though, the opposite is typically assumed—opinions are homogenous, but values are heterogeneous. While some might question the move against convention, it is quite sensible in situations where opinions diverge, but group members are all held accountable for any collective decisions made, e.g., policymaking. Thanks to an anonymous referee for suggesting that I flag this subtle but important point.

(**v**) No Regrets. For all admissible opinion profiles $(p_1,\ldots,p_n)$, $X \in \mathcal{A}$, and individuals i, $r_i^-(F(p_1,\ldots,p_n)^-(X)) = 0$ and $r_i^+(F(p_1,\ldots,p_n)^+(X)) = 0$ (in IP: for all admissible opinion profiles $(P_1,\ldots,P_n)$, $X \in \mathcal{A}$, and individuals i, $R_i^-(\mathcal{F}(P_1,\ldots,P_n)^-(X)) = 0$ and $R_i^+(\mathcal{F}(P_1,\ldots,P_n)^+(X)) = 0$).

The pragmatic criterion for pooling functions is seemingly intuitive, especially under the assumption that the individuals have a preference for the group not making bets that are regretted, *ex ante*. But even so, some might worry that satisfying the criterion will sometimes lead to abstaining from betting, and abstaining may cause group members to regret not betting. However, a regret of not betting will only manifest *ex post* due to thinking about what could have been obtained after learning an outcome that would have settled a bet the group considered making. That regret is different in kind, though, for it is not realized by any individual's regret functions, as defined above, and must be distinct from the regret that is of concern. With a focus on the specified *ex ante* regret only, the pragmatic criterion for pooling functions should come across as feasible.

2.3. *Some Observations*

I have so far introduced a few desirable criteria for pooling functions. Any pooling function satisfying the criteria should be considered plausible based on the epistemic and practical features inherited. With a concern only for F_L and $\mathcal{F}_A$, let us see which of the criteria each meets.

Observation 1. The linear pooling function, F_L, with domain Δ^n satisfies (**i**) Unanimity and (**ii**) Eventwise Independence.

Observation 1.1. The agglomerative pooling function, $\mathcal{F}_A$, with domain $\mathcal{D}^n$ satisfies (**i**) Unanimity.

Aczél and Wagner (1980) and McConway (1981) proved that F_L satisfies (**i**) and (**ii**). As for $\mathcal{F}_A$, when P_i, for all i, are identical, $\text{conv}(\cup_i P_i) = P_i$. Thus, the pooling function, $\mathcal{F}_A$, satisfies (**i**) Unanimity. While I do not have a proof showing that $\mathcal{F}_A$ satisfies (**ii**) Eventwise Independence, Stewart and Quintana (2018) proved that a (convex) agglomerative pooling function satisfies the criterion (see Proposition 2). Their result can be extended to our setting, but I leave that to the reader.

At this point, it appears that linear pooling and opinion agglomeration are both plausible strategies for aggregating opinions based on the above observations. But what about the other criteria?

Observation 2. The linear pooling function, F_L, with domain Δ^n satisfies (**iii**) Boundedness.

Observation 2.1. The agglomerative pooling function, $\mathcal{F}_A$, with domain $\mathcal{D}^n$ satisfies (**iii**) Boundedness and (**iv**) Non-Dictatorship.

(**iii**) Boundedness is a necessary property of F_L since any weighted average yielded by the function is in the closed interval $[\min(p_1(X),\ldots,p_n(X)),$ $\max(p_1(X),\ldots,p_n(X))]$, which is the set of all weighted averages of opinions, for all opinion profiles in the domain and $X \in \mathcal{A}$. (**iii**) Boundedness is a necessary property of $\mathcal{F}_A$ since any output of the function is a subset of the closed interval $\left[\inf \cup_i P_i(X), \sup \cup_i P_i(X)\right]$ by definition, for all opinion profiles in the domain and $X \in \mathcal{A}$. Furthermore, $\mathcal{F}_A$ satisfies (**iv**) Non-Dictatorship since every individual's opinion is contained in the output of the function. That is, $P_i(X) \subseteq \mathcal{F}_A(P_1,\ldots,P_n)(X)$, for all opinion profiles in the domain, individuals i, and $X \in \mathcal{A}$. Whether F_L satisfies (**iv**) Non-Dictatorship depends on the weights assigned to individuals. In case some *i* is given weight $w_i = 1$, F_L is dictatorial. Otherwise, if $0 < w_i < 1$, for all *i*, then F_L satisfies (**iv**) Non-Dictatorship. It may be stipulated then that no individual is given maximum weight to prevent pooling by F_L from collapsing into a dictatorship.[12]

With the added stipulation, both F_L and $\mathcal{F}_A$ again appear to be plausible aggregation strategies. So now, we only have one criterion left to consider, and the one that interests me the most in this paper.

Observation 3. The linear pooling function, F_L, with domain Δ^n violates (**v**) No Regrets.

Observation 3.1 The agglomerative pooling function, $\mathcal{F}_A$, with domain $\mathcal{D}^n$ satisfies (**v**) No Regrets.

Observation 3 is easy to see. Suppose that for some *non-unanimous* opinion profile in the domain and event $X \in \mathcal{A}$, $F_L(p_1,\ldots,p_n)^-(X) = F_L(p_1,\ldots,p_n)^+(X) =$

12. Although the pooling functions F_L and $\mathcal{F}_A$ can satisfy the usual criteria, there is one that both violate, namely, Probabilistic Independence. Laddaga (1977) highlighted that an agreement about the independence of events *X* and *Y* is not preserved by taking any non-extreme weighted average of opinions as a group's opinion, and Lehrer and Wagner (1983) observed that only dictatorial linear pooling functions satisfy Probabilistic Independence. Thus, F_L violates Probabilistic Independence when $0 < w_i < 1$ for all *i*. But in defense of non-dictatorial linear pooling, Lehrer and Wagner contend that individuals and groups may have very little concern for newly established probabilistic correlations between events since such dependencies are often artificial, resulting from the formal machinery. However, others have pointed out that the practical consequences are not negligible (see, e.g., Kyburg & Pittarelli 1996; Seidenfeld, Schervish, & Kadane 2010; Elkin & Wheeler 2018). As for $\mathcal{F}_A$, the resulting set-based opinions contain all convex combinations, so there is some $p \in \mathrm{conv}(\cup_i P_i)$ such that $p(X \mid Y) \neq p(X)$ for some admissible opinion profile and events *X* and *Y*. The reader should note, though, that independence is more complex in IP (see, e.g., Cozman 2012; Pedersen & Wheeler 2014).

$F_L(p_1,\dots,p_n)(X)$, but there is at least one j such that $p_j(X)=y$, $y\in\left[\min\left(p_1(X),\dots,p_n(X)\right),\max\left(p_1(X),\dots,p_n(X)\right)\right]$, and $y<F_L(p_1,\dots,p_n)(X)$. Then, $(F_L(p_1,\dots,p_n)^-(X)-y)>0$. Since $r_i^-(u)=\max\{u-p_i(X),0\}$, for all $i,u\in[0,\infty)$, and $X\in\mathcal{A}$, $r_j^-(F_L(p_1,\dots,p_n)^-(X))>0$. Therefore, F_L violates (**v**) No Regrets.

Observation 3.1 is also easy to see. For all admissible opinion profiles and events $X\in\mathcal{A}$, $\mathcal{F}_A(P_1,\dots,P_n)^-(X)=\inf\operatorname{conv}\left(\cup_i P_i(X)\right)$ and $\mathcal{F}_A(P_1,\dots,P_n)^+(X)=\sup\operatorname{conv}\left(\cup_i P_i(X)\right)$ by definition provided that $\mathcal{F}_A(P_1,\dots,P_n)(X)=\operatorname{conv}\left(\cup_i P_i(X)\right)$. Given the collective betting prices associated with the pooling function $\mathcal{F}_A$, it follows that $\max\left\{\mathcal{F}_A(P_1,\dots,P_n)^-(X)-P_i^-(X),0\right\}=0$ since $\left(\mathcal{F}_A(P_1,\dots,P_n)^-(X)-P_i^-(X)\right)\le 0$ and $\max\left\{P_i^+(X)-\mathcal{F}_A(P_1,\dots,P_n)^+(X),0\right\}=0$ since $(P_i^+(X)-\mathcal{F}_A(P_1,\dots,P_n)^+(X))\le 0$, for all admissible profiles $(P_1,\dots,P_n)$, events $X\in\mathcal{A}$, and individuals i. Therefore, $\mathcal{F}_A$ satisfies (**v**) No Regrets.

Observation 3 does not bode well for linear pooling. Those who favor it, however, might suggest weakening the No Regrets criterion by requiring instead that the *ex ante* regret is zero for *at least* one transaction type for all individuals i. Weakening the requirement may not only be an easy fix for linear pooling, but weakening could also be justified on the grounds that some might care more about avoiding buyer's remorse and consequently value zero *ex ante* regret with respect to r_i^- more than r_i^+ (or *vice versa* if loss averse). Unfortunately, weakening No Regrets in the way described leads to another problem for linear pooling.

> *Observation 4*. For all admissible opinion profiles $(p_1,\dots,p_n)$, events $X\in\mathcal{A}$, and individuals i, if $p_j(X)\neq p_k(X)$, for all j and k, $F_L(p_1,\dots,p_n)(X)=x$, and $r_i^-(x)=\max\{x-p_i(X),0\}=0$ or $r_i^+(x)=\max\{p_i(X)-x,0\}=0$, then F_L violates (**iv**) Non-Dictatorship.

Suppose that $p_j(X)\neq p_k(X)$, for all j and k, $F_L(p_1,\dots,p_n)(X)=x$, and $r_i^-(x)=\max\{x-p_i(X),0\}=0$. If $0<w_i<1$ for all individuals i, then x is in the open interval $\left(\min\left(p_1(X),\dots,p_n(X)\right),\max\left(p_1(X),\dots,p_n(X)\right)\right)$. But because $F_L(p_1,\dots,p_n)^-(X)=x$ and $r_i^-(x)=\max\{x-p_i(X),0\}=0$, for all i, then $x=\min\left(p_1(X),\dots,p_n(X)\right)$, which is not contained in $(\min\left(p_1(X),\dots,p_n(X)\right)$, $\max\left(p_1(X),\dots,p_n(X)\right))$. So it is not the case that $w_i<1$ for all individuals i, and there must be some individual i whose opinion $p_i(X)$ realizes $\min\left(p_1(X),\dots,p_n(X)\right)$ and is given weight $w_i=1$, but then, F_L violates (**iv**) Non-Dictatorship. Similarly, one can follow this chain of reasoning with the assumption that $r_i^+(x)=\max\{p_i(X)-x,0\}=0$ instead and arrive at the same conclusion. Furthermore, it is not clear that weakening No Regrets is a feasible fix. Notice that, under the above conditions, eliminating the perceived group loss on one side can significantly increase the perceived group loss on the other side from the viewpoints of some individuals when pooling by F_L.

In summary, linear pooling fails to satisfy all the desirable criteria (**i**)–(**v**), whereas opinion agglomeration satisfies all five. Note, however, that linear pooling satisfies (**i**)–(**v**) if $(p_1,\ldots,p_n)$ is *unanimous*, but in such instances, pooling is not of much interest since it would be trivial. As it turns out then, opinion agglomeration is the more plausible aggregation strategy of the two.

3. Justifying the Behavioral Criterion

Some might think that the reason I introduced No Regrets is to strategically undermine linear pooling. In this section, I will offer justification for the criterion, showing that it is indeed a reasonable constraint normatively and empirically to impose on opinion aggregation.

3.1 Regret as an Emotion and Its Avoidance

There has been much talk about regret in this paper, but what is it? In short, regret is considered a negative emotion in response to counterfactual reasoning about a fault in personal action (Roese & Summerville 2005), or simply, *counterfactual emotion* (Kahneman & Miller 1986). Described in more depth, Zeelenberg (1999) offers the following account.

> Regret is a cognitive emotion: it is an emotion that needs cognition to be experienced and that may produce cognitions as well. In order to *feel* regret one has to *think*. One has to think about one's choices and the outcomes generated by these choices, but one also has to think about what other outcomes might have been obtained by making a different choice. Thus, regret is typically felt in response to decisions that produce unfavorable outcomes compared to the outcomes that the rejected option would have produced. (1999: 327)

To give an example of a regrettable decision, consider the following.

> *Lemon or Peach*. You intend to purchase a used car today. You visit a local car dealership, hoping to leave with a vehicle. There, you find an attractive, used, mid-sized sedan priced at the top of your budget. You are familiar with the model, but you are uncertain whether it is of low quality, a lemon, or of high quality, a peach, in its used condition. You ultimately decide to buy the car and drive away a happy customer. But your happiness is short-lived, for when you arrive home, your neighbor, who

> is a professional mechanic, inspects the car and discovers a major fault with the engine. He says, "This vehicle will not pass a certified inspection and cannot be driven on public roads." Disappointment hits you, as you have come to learn that the car is a lemon.

Your troubles in the given scenario begin with an asymmetry in information.[13] As the buyer, you are uncertain whether the used car is a lemon or a peach, but the dealer knows which it is, and of course, will not reveal if it is a lemon. Despite the information asymmetry, you buy the car, but later learn that it is a lemon. Thinking about how the dealer successfully swindled the maximum from you for a vehicle they knew to be worth very little, you realize that it would have been wise to pass on it instead. That realization and accompanying emotional pain is your regret. While fictitious, the example mirrors countless experiences of consumer regret, which is a common side-effect of acting under uncertainty, and something that rational individuals should generally seek to avoid.

The *Lemon or Peach* case was introduced for the purpose of generating thoughts about what it is like to be in a distressing state of regret after a decision has been made and the true state of the world is learned, but to reiterate, the No Regrets criterion concerns *ex ante* regret. Since the latter kind of regret is thought about less often, though, some might wonder what can cause it. Consider, for example, the collective behavior in a variation of *Lemon or Peach*.

> *Lemon or Peach–Group*. You and your partner intend to purchase a used car today with a maximum budget of $10,000. You go together to a local car dealership, hoping to find a vehicle. There, an attractive, used, mid-sized sedan priced at $10,000 catches your attention. You and your partner are both familiar with the model but are uncertain of what it is worth in its used condition. You say, "I think it is worth the asking price, so we should buy it." Your partner, however, says, "I think that it is worth no more than $6,000." In light of this dispute, you suggest compromising and offering the dealer $8,000. Your partner is still reluctant to believe that the car is worth more than $6,000, but they agree, *regretfully*, to making the bid for the sake of compromising. The dealer accepts, and you jointly purchase the car. When you arrive home, your neighbor, who is a car appraiser at a prestigious appraisal firm, yells over at the two of you in a sarcastic tone, "How much did you pay for that thing?" You both reply, "$8,000." She replies, "In that condition, the car is worth at most $6,000!"

13. The *Lemon or Peach* case is inspired by the main example of George Akerlof's (1970) seminal paper, "The Market for Lemons: Quality, Uncertainty, and the Market Mechanism."

A notable difference here from the first case, aside from an individual/group disparity, is that an additional state of regret is experienced, which occurs before learning the actual value of the car. The bid itself was immediately regretted by your partner. That regret is thus *ex ante* rather than *ex post*.

A lesson that can be drawn from the second case, and that further supports a principle like No Regrets, is that if you had heeded the warning of your partner, the two of you may have avoided the distress later felt. Put another way, had the two of you compromised in a fashion such that neither of you regretted the group's bidding price, the group would have decided differently, resulting in less grief. So the No Regrets criterion is not only justified on the basis of preventing group members from experiencing the unwanted emotion, *ex ante*, but also on the grounds that it can be instrumental in preventing group members from experiencing regret, *ex post*. For an expressed *ex ante* regret by a group member forewarns of an *ex post* regret that is still to come, just as in the second case.

3.2. *Aversion to Making Bad Deals: Experimental Evidence*

It is widely known that the judgments and decisions of actual, not idealized, individuals often deviate from those predicted by expected utility theory. When endowed with a good, for instance, an average person will likely be reluctant to sell the good for the same price that they would buy it, giving rise to the so-called *endowment effect* (Thaler 1980). Tversky and Kahneman's (1991) explanation for why the endowment effect arises is that utility is reference-dependent, exchanging a possession for a modest gain is viewed as a loss relative to an individual's reference point, and individuals are averse to losses.

For decades, many behavioral economists have attributed the endowment effect to an aversion to losing possessions. However, recent evidence casts doubt on the view that loss aversion is the only plausible explanation for a gap between 'willingness to pay' and 'willingness to accept'. Weaver and Frederick (2012) hypothesized that a difference between buying and selling prices for a good is sometimes a result of an aversion to making bad deals. They predicted that by manipulating a subject's reference price for a good, 'willingness to pay' would largely differ from 'willingness to accept'. If the prediction is true, then the endowment effect is a consequence of an aversion to making bad deals as opposed to an aversion to losing the good, for the relevant factor is the reference price, not the good itself. The results of their experiments confirmed the prediction.

In their first experiment, for example, participants were presented with four different kinds of candy typically sold at movie theaters and asked to indicate their favorite. Participants were either endowed or not endowed with their choice and assigned to a high or moderate reference price condition. In the high condition, participants were informed that their chosen box of candy sells for \$4.00 at the Harvard Square Theater (near Harvard University). In the moderate condition, participants were told that the box of candy sells for \$1.49 at a local Target retail store. The experimenters then elicited the maximum buying and minimum selling prices from participants following the Becker, DeGroot, and Marschak (1964) method. They found in the high condition a significant difference in prices ($\mu_{sell} = \$2.88$, $\mu_{buy} = \$1.54$) (Weaver & Frederick 2012: 698). Although these results only illustrate a higher average selling price upon inflating participants' reference prices, I refer the reader to the results of Experiments 3a and 3b, which confirm a lower average buying price upon depressing participants' reference prices.

What can we make of these experimental results? A first observation is that individuals appear to be keen on avoiding *ex ante* regret when transacting due to an aversion to making bad deals. Whether groups exhibit similar behavior is an open question, but the No Regrets criterion is at least empirically tenable, as it quantifies over individuals in a group. A second observation is that individuals appear to be inclined to revise their judgments after receiving information that shifts their reference prices, raising the question: will individuals form imprecise opinions after learning the opinions of other group members due to an aversion to making bad betting deals? Although I am unable to pursue the question in this paper, the empirical results discussed have interesting implications for a behavioral theory of individual and group opinions and pave the way for future work.

4. Conclusion

In summary, I showed that opinion agglomeration fares better than linear pooling in satisfying the desirable criteria laid out in Section 2. Trouble for linear pooling surfaced when considering the pragmatic criterion, No Regrets, since aside from pooling a unanimous opinion profile, linear pooling fails to satisfy it. Even under a weakened version, linear pooling was shown to be no better off, as it then violates the Non-Dictatorship criterion. After presenting the formal results, I proceeded to give normative and empirical arguments in support of the pragmatic criterion for opinion aggregation, and that ultimately favor opinion agglomeration.

Acknowledgments

I owe many thanks to Stephan Hartmann, Richard Pettigrew, Tomasz Żuradzki, Mikaël Cozic, Ophelia Deroy, Bahador Bahrami, and two anonymous referees for their valuable comments that tremendously improved the ideas in this paper. I would also like to thank audiences at Jagiellonian University in Kraków, Ludwig Maximilians Universität München, Tilburg University, and the "Reasoning in Social Context" colloquium at The Royal Netherlands Academy of Arts and Sciences. The research for this article was generously supported by the Agence Nationale de la Recherche (ANR) (grant No. ANR-16-FRAL-0010), the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant No. 805498), and the NOMIS Foundation (project "Diversity in Social Environments").

References

Aczél, János and Carl Wagner (1980). A Characterization of Weighted Arithmetic Means. *SIAM Journal on Algebraic Discrete Methods*, *1*(3), 259–60.

Akerlof, George A. (1970). The Market for "Lemons": Quality, Uncertainty, and the Market Mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500.

Augustin, Thomas, Frank P. A. Coolen, Gert De Cooman, and Matthias C. M. Troffaes (2014). *Introduction to Imprecise Probabilities*. John Wiley and Sons.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964). Measuring Utility by a Single-Response Sequential Method. *Behavioral Science*, *9*(3), 226–32.

Bradley, Seamus (2019). Imprecise Probabilities. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Retrieved from https://plato.stanford.edu/entries/imprecise-probabilities

Christensen, David (2007). Epistemology of Disagreement: The Good News. *The Philosophical Review*, *116*(2), 187–218.

Christensen, David (2011). Disagreement, Question-begging, and Epistemic Self-Criticism. *Philosophers' Imprint*, *11*(6), 1–22.

Couso, Inés and Serafín Moral (2009). Sets of Desirable Gambles and Credal Sets. In *Proceedings of the 6th International Symposium on Imprecise Probability: Theories and Applications* (99–108).

Cozman, Fabio G. (2012). Sets of Probability Distributions, Independence, and Convexity. *Synthese*, *186*(2), 577–600.

DeGroot, Morris H. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, *69*(345), 118–21.

Dietrich, Franz (2019). A Theory of Bayesian Groups. *Noûs*, *53*(3), 708–36.

Dietrich, Franz and Christian List (2016). Probabilistic Opinion Pooling. In Alan Hájek and Christopher Hitchcock (Eds.), *The Oxford Handbook of Philosophy and Probability* (519–44). Oxford University Press.

Dubois, Didier and Henri Prade (1988). *Possibility Theory: An Approach to Computerized Processing of Uncertainty*. Plenum.

Easwaran, Kenny, Luke Fenton-Glynn, Christopher Hitchcock, and Joel D. Velasco (2016). Updating on the Credences of Others: Disagreement, Agreement, and Synergy. *Philosophers' Imprint*, *16*(11), 1–39.

Elga, Adam (2007). Reflection and Disagreement. *Noûs*, *41*(3), 478–502.

Elkin, Lee and Gregory Wheeler (2018). Resolving Peer Disagreements through Imprecise Probabilities. *Noûs*, *52*(2), 260–78.

Ellsberg, Daniel (1961). Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics*, *75*(4), 643–69.

Feldman, Richard (2005). Respecting the Evidence. *Philosophical Perspectives*, *19*(1), 95–119.

Feldman, Richard (2007). Reasonable Religious Disagreements. In Louise Antony (Ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life* (194–214). Oxford University Press.

Finetti, Bruno de (1974). *Theory of Probability: A Critical Introductory Treatment*. Wiley.

Galton, Francis (1907). Vox Populi (The Wisdom of Crowds). *Nature*, *75*(1949), 450–51.

Garg, Ashutosh, Thathachar S. Jayram, Shivakumar Vaithyanathan, and Huaiyu Zhu (2004). Generalized Opinion Pooling. In *Proceedings of the 8th International Symposium on Artificial Intelligence and Mathematics* (79–86).

Genest, Christian and James V. Zidek (1986). Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, *1*(1), 114–35.

Gilboa, Itzhak and Massimo Marinacci (2016). Ambiguity and the Bayesian Paradigm. In Horacio Arló-Costa, Vincent F. Hendricks, and Johan van Benthem (Eds.), *Readings in Formal Epistemology* (385–439). Springer.

Gilboa, Itzhak and David Schmeidler (1989). Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics*, *18*(2), 141–53.

Hartmann, Stephan (2017). Prospect Theory and the Wisdom of the Inner Crowd. Unpublished manuscript.

Hooker, Reginald H. (1907). Mean or Median. *Nature*, *75*(1951), 487.

Jehle, David and Branden Fitelson (2009). What Is the "Equal Weight View"? *Episteme*, *6*(3), 280–93.

Joyce, James M. (2005). How Probabilities Reflect Evidence. *Philosophical Perspectives*, *19*(1), 153–78.

Joyce, James M. (2010). A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives*, *24*(1), 281–323.

Kahneman, Daniel and Dale T. Miller (1986). Norm Theory: Comparing Reality to its Alternatives. *Psychological Review*, *93*(2), 136–53.

Keller, L. Robin, Rakesh K. Sarin, and Jayavel Sounderpandian (2007). An Examination of Ambiguity Aversion: Are Two Heads Better than One? *Judgment and Decision Making*, *2*(6), 390–97.

Kelly, Thomas (2013). Disagreement and the Burdens of Judgment. In David Christensen and Jennifer Lackey (Eds.), *The Epistemology of Disagreement: New Essays* (31–53). Oxford University Press.

Konek, Jason (in press). Epistemic Conservativity and Imprecise Credence. *Philosophy and Phenomenological Research*.

Kyburg, Henry E. and Michael Pittarelli (1996). Set-Based Bayesianism. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, *26*(3), 324–39.

Laddaga, Robert (1977). Lehrer and the Consensus Proposal. *Synthese*, *36*(4), 473–77.

Lehrer, Keith (1976). When Rational Disagreement is Impossible. *Noûs*, *10*(3), 327–32.

Lehrer, Keith and Carl Wagner (1983). Probability Amalgamation and the Independence Issue: A Reply to Laddaga. *Synthese*, *55*(3), 339–46.

Levi, Isaac (1974). On Indeterminate Probabilities. *The Journal of Philosophy*, *71*(13), 391–418.

Levi, Isaac (1983). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press.

Loomes, Graham and Robert Sugden (1982). Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *The Economic Journal*, *92*(368), 805–24.

Lyon, Aidan (2017). Vague Credences. *Synthese*, *194*(10), 3931–54.

Machina, Mark J. and Marciano Siniscalchi (2014). Ambiguity and Ambiguity Aversion. In Mark J. Machina and W. Kip Viscusi (Eds.), *Handbook of the Economics of Risk and Uncertainty* (Vol. 1, 729–807). North-Holland.

McConway, Kevin J. (1981). Marginalization and Linear Opinion Pools. *Journal of the American Statistical Association*, *76*(374), 410–14.

Moral, Serafín and José del Sagrado (1998). Aggregation of Imprecise Probabilities. In Bernadette Bouchon-Meunier (Ed.), *Aggregation and Fusion of Imperfect Information: Studies in Fuzziness and Soft Computing* (Vol. 12, 162–88). Physica.

Moss, Sarah (2011). Scoring Rules and Epistemic Compromise. *Mind*, *120*(480), 1053–69.

Nau, Robert F. (2002). The Aggregation of Imprecise Probabilities. *Journal of Statistical Planning and Inference*, *105*(1), 265–82.

Pedersen, Arthur P. and Gregory Wheeler (2014). Demystifying Dilation. *Erkenntnis*, *79*(6), 1305–42.

Pettigrew, Richard (2019). On the Accuracy of Group Credences. In Tamar Szabó Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 6, 137–60). Oxford University Press.

Roese, Neal J. and Amy Summerville (2005). What We Regret Most . . . and Why. *Personality and Social Psychology Bulletin*, *31*(9), 1273–85.

Russell, Jeffrey S., John Hawthorne, and Lara Buchak (2015). Groupthink. *Philosophical Studies*, *172*(5), 1287–1309.

Savage, Leonard J. (1951). The Theory of Statistical Decision. *Journal of the American Statistical Association*, *46*(253), 55–67.

Savage, Leonard J. (1954). *The Foundations of Statistics*. Dover.

Schoenfield, Miriam (2012). Chilling Out on Epistemic Rationality. *Philosophical Studies*, *158*(2), 197–219.

Seidenfeld, Teddy, Mark J. Schervish, and Joseph B. Kadane (2010). Coherent Choice Functions under Uncertainty. *Synthese*, *172*(1), 157–76.

Shafer, Glenn (1976). *A Mathematical Theory of Evidence*. Princeton University Press.

Simon, Herbert A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.

Stewart, Rush T. and Ignacio O. Quintana (2018). Probabilistic Opinion Pooling with Imprecise Probabilities. *Journal of Philosophical Logic*, *47*(1), 17–45.

Stone, Mervyn (1961). The Opinion Pool. *The Annals of Mathematical Statistics*, *32*(4), 1339–42.

Sturgeon, Scott (2010). Confidence and Coarse-Grained Attitudes. In Tamar Szabó Gendler and John Hawthorne (Eds.), *Oxford Studies in Epistemology* (Vol. 3, 126–49). Oxford University Press.

Surowiecki, James (2005). *The Wisdom of Crowds*. Anchor.

Thaler, Richard (1980). Toward a Positive Theory of Consumer Choice. *Journal of Economic Behavior & Organization*, *1*(1), 39–60.

Tversky, Amos and Daniel Kahneman (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. *The Quarterly Journal of Economics*, *106*(4), 1039–61.

Wagner, Carl (1978). Consensus through Respect: A Model of Rational Group Decision-making. *Philosophical Studies*, *34*(4), 335–49.

Wagner, Carl (2009). Jeffrey Conditioning and External Bayesianity. *Logic Journal of IGPL*, *18*(2), 336–45.

Walley, Peter (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.

Wasserman, Larry A. and Joseph B. Kadane (1990). Bayes' Theorem for Choquet Capacities. *The Annals of Statistics*, *18*(3), 1328–39.

Weaver, Ray and Shane Frederick (2012). A Reference Price Theory of the Endowment Effect. *Journal of Marketing Research*, *49*(5), 696–707.

Wheeler, Gregory (2020). Bounded Rationality. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Retrieved from https://plato.stanford.edu/entries/bounded-rationality/

Zeelenberg, Marcel (1999). Anticipated Regret, Expected Feedback, and Behavioral Decision Making. *Journal of Behavioral Decision Making*, *12*(2), 93–106.