Scoring rules and epistemic compromise

Sarah Moss
moss@mit.edu

Formal models of epistemic compromise have several fundamental applications. Disagreeing agents may construct a compromise of their opinions to guide their collective action. Or they may construct a compromise to guide them in updating their individual credences. Recent literature on disagreement has focused on certain questions about epistemic compromise: when you find yourself disagreeing with an epistemic peer, under what circumstances, and to what degree, should you change your credence in the disputed proposition? Elga 2006 and Christensen 2007 say you should compromise often and deeply; Kelly 2007 disagrees. But these authors leave open another question: what constitutes a perfect compromise of opinion?

In the disagreement literature, it is sometimes assumed that if we assign different credences c_a and c_b to a proposition p, we reach a perfect compromise by *splitting the difference* in our credences. In other words: to adopt a perfect compromise of our opinions is to assign credence $.5(c_a + c_b)$ to p. For instance, Kelly 2007 says that when peers assign .5 and .7 to a proposition, to adopt a compromise is to "split the difference with one's peer and believe the hypothesis to degree .6" (19).[1]

But why does .6 constitute a perfect compromise? Of course, .6 is the arithmetic mean of .5 and .7. But why must a compromise of agents' opinions always be the arithmetic mean of their prior credences? In other cases of compromise, we do not simply take it for granted that the outcome that

1. See Christensen 2007, Joyce 2007, and Weatherson 2007 for further recent references to compromising by splitting the difference.

constitutes a perfect compromise is determined by the arithmetic mean of quantities that reflect what individual agents most prefer. Suppose we are running partners, and I want to run one mile, while you want to run seven. Running four miles may not be a perfect compromise, especially if I strongly prefer running one mile over running four, while you only slightly prefer running farther.

The moral of this paper is that the same sort of situation may arise in purely epistemic cases of compromise, cases in which each agent prefers having certain credences over others, where this preference is grounded in purely epistemic concerns. Suppose I strongly prefer assigning .1 credence to a disputed proposition, while you weakly prefer credences around .7. In this kind of case, we may reach a perfect compromise by converging on some shared credence *lower than* .4. Splitting the difference may not constitute a perfect compromise when agents who have different credences also have different epistemic values.

To make this moral precise, we must say how an agent may value certain credences over others, in a purely epistemic sense. I take an agent's *scoring rule* to measure how much she epistemically values various alternative credences she might assign a given proposition. It is natural to suppose that agents should assess just this kind of value as they judge how much they would prefer certain consensus opinions over others. Using scoring rules, we can develop a natural alternative to the strategy of compromising by splitting the difference: agents may compromise by coordinating on the credences that they collectively most prefer, given their epistemic values.

I have two main aims in this paper: to develop this alternative strategy, and to argue that this strategy governs how agents should compromise. In §1, I define the notion of a scoring rule and introduce relevant properties of these epistemic value functions. In §2, I develop the alternative strategy for compromising that I defend. In §3, I compare my alternative strategy with the traditional strategy of splitting the difference. I characterize the situations in which the two strategies coincide, and those in which they differ. In §4, I argue that where the strategies do yield different recommendations, compromising by maximizing epistemic value is a reasonable strategy. Finally, in §5, I discuss applications of the strategy I defend.

1 Scoring rules

In assigning a particular credence to a proposition, you are estimating its truth value. In a purely epistemic sense, closer estimates are better. Having .8 credence in a truth is better than having .7 credence, and having .9 credence is better still.[2]

How much better? Different agents may value closer estimates differently. For instance, suppose you care a lot about your credences as they approach certainty. In this case, you may value having .9 credence in a truth much more than having .8 credence, without valuing .6 much more than .5. Or suppose you do not particularly care about having credences that approach certainty, as long as your beliefs are on the right track. In that case, you may equally value .9 and .8 credence in a truth, and equally value .2 and .3 credence in a falsehood.

Facts like these are traditionally modelled by your *scoring rule*, a record of how much you value various estimates of truth values.[3] Formally, a scoring rule f is a pair of functions f_1, f_0 from $[0,1]$ to $\mathbb{R}$. Intuitively, the first function, $f_1(x)$, measures how much you value having credence x in a proposition that turns out to be true. For instance, if you value having .9 credence in a true proposition much more than having .8 credence, without valuing .6 much more than .5, then the first function f_1 of your scoring rule will reflect this preference:

$$[f_1(.9) - f_1(.8)] > [f_1(.6) - f_1(.5)].$$

The second function $f_0(x)$ measures the value of having credence x in a proposition that turns out to be false. For instance, if it makes no difference whether you have .2 or .3 credence in a falsehood, then the second function

2. For further discussion of the notion of credences as estimates, see Jeffrey 1986 and Joyce 1998. For further discussion of the notion of *purely epistemic* value, compare the value of having accurate credences with the more familiar value of having true beliefs, as discussed in Alston 2005 and Lynch forthcoming.

3. Scoring rules were independently developed by Brier 1950 and Good 1952 to measure the accuracy of probabilistic weather forecasts. Savage 1971 uses scoring rules to assess forecasts of random variables, and treats assignments of credence to particular events as a special case of such forecasts. For more recent literature using scoring rules to assess credences, see Gibbard 2006, Joyce 1998, and Percival 2002.

f_0 of your scoring rule will have equal values on these credences: $f_0(.2) = f_0(.3)$.

Now we can formally model the statement that closer estimates of truth value are better, i.e. more valuable to you. Closer estimates of the truth value of propositions that turn out to be true are more valuable to you just in case your scoring rule $f_1(x)$ is strictly increasing. Closer estimates of the truth value of falsehoods are more valuable just in case $f_0(x)$ is strictly decreasing.

Taken together, your scoring rule and your actual credence in a proposition p determine the *expected epistemic value* of your having a particular credence in p. Suppose you actually have credence α in p. Then your expected epistemic value of having some other credence x in p is just a weighted sum of the value you assign to having credence x in p if p is true, and to having credence x in p if p is false:

$$\textsc{ev}(x, \alpha, f_1, f_0) \quad =_{df} \quad \alpha f_1(x) + (1 - \alpha) f_0(x)$$

Your scoring rule says how much value you assign to various credences, including those other than your own. Note that even if you and another agent have the same scoring rule, you may well assign your own credence in p a greater expected epistemic value than you assign her credence, given your own assessment of how likely it is that p is true.

Moreover, you may even assign your own credence in p a greater expected value than you assign *any* alternative credence. If this is always the case, then your scoring rule is *credence-eliciting*. In other words, your scoring rule is credence-eliciting when no matter what credence you assign to a proposition p, assigning that credence maximizes the expected value of your credence in p. In other words, you always do the best you can by your own lights by assigning p the credence that you do.

Using your scoring rule and your current credence in a proposition, we can calculate the expected value of your having various alternative credences in that proposition. Before applying scoring rules to cases of compromise, I should define one more useful measure: the expected value of your having a particular *credence distribution* over an algebra of many propositions. It is relatively simple to come up with a natural measure of expected value for credence distributions. Roughly, we can compute the expected value of your having several credences in separate propositions by simply summing up the expected value of your having each of those individual credences.

Formally, let Γ be your actual credence distribution over some algebra P, and let Q be the set of the *atomic propositions* of P, i.e. the strongest propositions in P that together form a partition of the space of all possible worlds. Let D be any prospective credence distribution over the P propositions. It is natural to define the expected value of having credence distribution D over P as the sum of the expected values of having the credences assigned by D to the atomic propositions of P:

$$\textsc{ev}(D, \Gamma, f_1, f_0) \quad =_{df} \quad \sum_{q_i \in Q} \textsc{ev}(D(q_i), \Gamma(q_i), f_1, f_0)$$

In other words, the expected value of a credence distribution is the sum of the expected values of the credences it assigns to atomic propositions.

To take a simple example, suppose you only have opinions about a single proposition p. Then your credence distribution is defined over an algebra with atomic propositions p and $\neg p$. The expected value of your having any alternative credence distribution D over this algebra is simply the sum of the expected value of your having credence $D(p)$ in p and credence $D(\neg p)$ in $\neg p$.

2 Compromise beyond splitting the difference

As long as your scoring rule is credence-eliciting, you prefer your credence in p to any other. That is, as long as your rule is credence-eliciting, your scoring rule will never dictate that you should change your credence in a proposition to another credence with greater expected epistemic value. It is widely accepted that this means that if you are rational, you must have a credence-eliciting scoring rule.[4] Otherwise, your scoring rule could motivate you to raise or lower your credences *ex nihilo*, in the absence of any new evidence whatsoever.

For this reason, how rational agents value alternative credences in a proposition is rarely of practical relevance. Of course, your scoring rule

4. See Oddie 1997, Joyce 1998 and Gibbard 2006 for further discussion.

affects how well you think other agents are doing, by affecting the expected value you assign to their credences. But if you are rational, you never value another credence more than your own. So your valuing of non-actual credences will never influence your behavior. For this reason, as long as your scoring rule is credence-eliciting, it will even be hard to tell exactly which scoring rule you have.

However, cases of compromise provide a practical application for the notion of a scoring rule. Scoring rules are practically relevant when agents have to assess the epistemic value of alternative credences in a proposition. Cases of compromise call for exactly this. Compromising agents must pick a credence to coordinate on. In doing so, they have to assess the epistemic value of alternative credences, if they are to determine which shared credence they most prefer.

This points the way to an alternative strategy for constructing a compromise of agents' opinions: *maximizing expected epistemic value*. Suppose we assign different credences to some proposition p, but we must construct a credence that constitutes the compromise of our opinions. The following seems like a natural strategy: choose our consensus credence in p to maximize the average of the expected values that we each assign to alternative credence distributions over the algebra with atomic propositions p and $\neg p$.[5]

For instance, suppose you have credence distribution A and I have credence distribution B, and suppose that you score credences with f_1 and f_0, and I score them with g_1 and g_0. Then we should compromise by choosing the credence distribution D that maximizes:

$$\textsc{aev}(D, A, f_1, f_0, B, g_1, g_0) \quad =_{df} \quad \frac{1}{2}[\textsc{ev}(D, A, f_1, f_0) + \textsc{ev}(D, B, g_1, g_0)]$$

5. In order to implement this strategy, we need only assess the epistemic value of credence distributions over simple four-element algebras; the equivalence results I give are restricted to compromises within this relevant class of credence distributions.

This is our alternative compromise strategy: rather than just splitting the difference between their credences, we may compromise by maximizing the average of the expected values we give to our consensus credence distribution.[6]

I will remain neutral on questions central to the disagreement literature, such as whether peers should trade in their prior credences for a perfect compromise of their opinions. This question is independent of questions about what constitutes a perfect compromise. For those engaged in the disagreement literature, deciding what constitutes a perfect compromise is just the beginning. For instance, those who think peers' opinions should converge on some consensus must also decide: once peers choose a consensus credence, how exactly should they revise their prior credence distributions to adopt that credence in the disputed proposition? In particular, how might their credences in non-disputed propositions change? I will set this question aside, though see Fitelson & Jehle 2007 for further discussion.

I will demonstrate that agents who compromise by maximizing epistemic value rarely end up simply splitting the difference in their credences. Suppose you meet agents with exactly the same credences, but with different scoring rules. If you were to compromise with each agent individually, you may end up constructing different consensus opinions with each of them, even if your scoring rule is credence-eliciting. This means that if an agent compromises with others by maximizing expected epistemic value, her behavior will carry a lot of information about her scoring rule. In other words, we have discovered that a way in which rational agents value alternative credences can be a matter of practical relevance. Hence cases of compromise give us valuable motivation for studying scoring rules: as I will argue, they are cases in which previously inert differences in scoring rules make for real differences in what compromising agents should do.

6. There are several properties we might wish for in a procedure for aggregating credences; it is notoriously difficult to find a procedure that is satisfactory in all respects. See French 1985 for a discussion of impossibility results, and Genest & Zidek 1986 for a canonical overview and bibliography of the consensus literature. Shogenji (2007) and Fitelson & Jehle (2007) draw on this literature to present challenges for the strategy of splitting the difference. I will set aside traditional debates about judgment aggregation in order to focus on the relevance of scoring rules to norms governing compromise.

3 Comparing compromise strategies

In some cases, our alternative strategy yields a traditional recommendation. In particular, when agents *share a credence-eliciting scoring rule*, they maximize the average expected value of prospective consensus credence distributions by simply splitting the difference between their credences. For example, suppose we value prospective credences using the *Brier score:*[7]

$$\begin{aligned} f_0(x) &= 1 - x^2 \\ f_1(x) &= 1 - (1-x)^2 \end{aligned}$$

Suppose you have .7 credence in p and I have .1 credence. The Brier score is a credence-eliciting scoring rule. Hence to maximize the average of our expected values for prospective shared credence distributions, we should compromise by giving .4 credence to p.

Sometimes agents pursue a perfect compromise, one that is as fair and even as possible. But agents may also pursue an imperfect compromise, one that favors some agents more than others. For example, when an expert and an amateur disagree, they may elect to compromise in a way that favors the expert's original opinion. One traditional strategy for generating an imperfect compromise is to take a weighted arithmetic mean of agents' original opinions.[8] For example, an amateur may give the prior opinion of an expert four times as much weight as his own. Then if the expert initially has .1 credence in p and the amateur has .7 credence, they may compromise at $.8(.1) + .2(.7) = .22$.

Our alternative compromise strategy can also be extended to cases of imperfect compromise. In a case of imperfect compromise, agents may maximize a weighted arithmetic mean of the expected values they give to their consensus credence distribution. Furthermore, the above equivalence result extends to cases of imperfect compromise. In the general case: when

---

7. Strictly speaking, traditional scoring rules measure the *inaccuracy* of a credence, and hence the *disvalue* of having that credence. For instance, Brier 1950 scores credences using the rule $f_0(x) = x^2$, $f_1(x) = (1-x)^2$. For simplicity, I follow Gibbard 2006 in using versions of scoring rules that measure positive epistemic value.

8. For instance, see the discussion of epistemic deference in Joyce 2007.

agents share a credence-eliciting rule, they maximize the weighted average of the expected values they give to their consensus credence distribution when the consensus is the *same weighted average* of their prior credences.[9] It follows that when agents share a credence-eliciting scoring rule, they maximize the exact average of their expected epistemic values by exactly splitting the difference between their prior credences.

Hence in some cases, including cases of imperfect compromise, splitting the difference and maximizing expected value coincide. But it is not hard to see that in principle, these strategies could yield different results. Sometimes agents who disagree about what is *practically* valuable do not maximize their expected utility by splitting the difference between their most preferred outcomes, as in the case where I want to run one mile and you want to run seven. Similarly, agents who disagree about what is epistemically valuable do not always maximize expected epistemic value by splitting the difference in their credences.

For instance, agents with different scoring rules may not maximize expected epistemic value by splitting the difference. Suppose you value prospective credences using the Brier score, and I value them using the following credence-eliciting rule:

$$\begin{aligned} g_0(x) &= x + \log(1-x) \\ g_1(x) &= x \end{aligned}$$

Suppose you have .7 credence in p and I have .1 credence. Even though our scoring rules are each credence-eliciting, our perfect compromise is asymmetric: in constructing a compromise, we maximize our expected epistemic value by choosing a consensus credence of approximately .393, not by splitting the difference.[10]

Even when agents share a scoring rule, they may not maximize expected epistemic value by splitting the difference in their credences. For example, suppose we both value prospective credences using the following

9. See Corollary of appendix for proof.

10. See Example 1 of appendix for details.

"Brier cubed" score:

$$\begin{aligned} h_0(x) &= 1 - x^3 \\ h_1(x) &= 1 - (1-x)^3 \end{aligned}$$

Suppose you have .7 credence in p and I have .1 credence. Then our perfect compromise is again asymmetric: in constructing a compromise, we maximize our expected epistemic value by choosing a consensus credence of approximately .449.[11]

This value has an interesting property: it is exactly what an agent using the Brier cubed score would assign to p to maximize the expected value of her new credence distribution, if she started with precisely .4 credence in p.[12] This is no coincidence: whenever agents share a scoring rule, they maximize the average of their expected values for prospective credence distributions by picking the distribution that has maximal expected value for an agent with their same scoring rule and the average of their credences. Furthermore, this claim extends again to a result about weighted averages of agents' expected epistemic values. That is, agents who share a scoring rule can maximize a given weighted arithmetic mean of their expected values for consensus credence distributions by following a straightforward rule; namely, they should choose the credence distribution that a hypothetical agent would most prefer, if she shared their scoring rule, and if her credence in p were just that same weighted arithmetic mean of their actual credences.[13]

This final result incorporates many results given so far. Compromising agents maximize the weighted average of their expected values by choosing the credence distribution preferred by an agent with the same weighted average of their credences. In the special case where compromising agents share a *credence-eliciting* rule, the hypothetical agent with that rule would most prefer her own credence. So agents sharing a credence-eliciting rule should compromise at the weighted average of their actual credence distributions. If the compromising agents pursue a perfect compromise, they

11. See Example 2 of appendix for details.

12. See Example 3 of appendix for details.

13. See Theorem of appendix for proof.

maximize their expected epistemic value by splitting the difference in their credences. On the other hand, we have seen that there are a number of cases in which our alternative compromise strategy comes apart from splitting the difference. It just remains to be shown that the former strategy can be more defensible than the latter.

4 Norms governing compromise

In order to understand how compromising agents should be influenced by their epistemic value functions, we must first understand how a single agent should be influenced by her own epistemic values. In the literature on scoring rules, several theorists have addressed the latter question. Our aim is to generalize their suggestions to norms governing multiple agents at once.

It is generally accepted that on pain of irrationality, a single agent must aim to maximize the expected epistemic value of her credences.[14] For example, Percival 2002 compares epistemic value with practical utility:

> Cognitive decision theory is a cognitive analog of practical decision theory... Bayesian decision theory holds that a rational agent...maximise[s] expected utility. Similarly, Bayesian cognitive decision theory holds that a rational cogniser...maximise[s] expected cognitive utility. (126)

In the same vein, Oddie 1997 says that scoring rules measure "a (pure) cognitive value which it is the aim of a rational agent, qua pure inquirer, to maximize" (535).

How can we extend this condition to a norm that applies to compromising agents? It is useful to first consider the following single agent case: suppose an evil scientist tells you he is going to perform an operation to change your credence in a certain proposition p, which you currently believe to degree .7. The scientist gives you a choice: after the operation, you may have either .6 or .8 credence in p. On pain of irrationality, you should choose whichever credence has the greater expected epistemic value. In general, if you are forced to adjust your credence in p, so that your new

14. Or at least she must maximize expected epistemic value, given that her epistemic values are themselves rationally permissible, e.g. credence-eliciting. See Percival 2002 for further discussion.

credence satisfies certain constraints, you should choose the alternative credence with the greatest expected epistemic value for you.

Cases of compromise are relevantly similar. In order to adopt a compromise of their prior opinions, agents are forced to adjust their credences in p so that they satisfy certain constraints. Only in cases of compromise, these constraints are defined extrinsically rather than intrinsically; namely, their adjusted credences must be equal to each other. In this situation, agents should choose the alternative credences with the greatest possible expected epistemic value for them.

In saying precisely how agents should construct a compromise of their opinions, it is useful to see that we face an analogous question when we extend norms governing expected practical utility to cases of practical compromise. Suppose we are deciding where to go for dinner. Based on my credences and practical utilities, I slightly prefer Chinese. Based on your credences and practical utilities, you strongly prefer Indian. Our natural response to the situation is that your stronger preference matters more: we should maximize the average of our expected utilities by choosing Indian. Roughly the same principle governs many reasonable voting systems. Every voter decides which outcome he thinks is most likely to make him the most satisfied. His vote reflects his expected utilities, and the election outcome reflects the average of many agents' expected utilities.

I do not mean to suggest that maximizing average expected utility is an ideal decision procedure. But in many situations, maximizing average expected utility is an intuitively reasonable method of deciding what to do when we have different practical utility functions and want to choose a maximally fair action. Epistemic value is the cognitive analog of practical utility. So we need good reason not to use the same method when deciding what to believe, when we have different epistemic utility functions and want to choose a maximally fair credence. In other words, we need good reason to avoid aggregating epistemic preferences in the way we generally aggregate practical preferences. It is not as if we must weigh practical and epistemic utilities in deciding what credence to assign: in epistemology contexts, epistemic values—encoded in scoring rules—are the only values at issue.

One additional reason to prefer the alternative strategy is that agents who compromise by maximizing the average of their expected epistemic values will never both prefer an alternative shared credence. Other compromise strategies—including most strategies that make no reference to epistemic value—could in principle lead agents to compromise on one credence when there is an alternative that both agents would independently prefer. And that kind of prescription would sit uncomfortably with the norm that every agent should independently aim to maximize her expected epistemic value. If epistemic values should influence an individual in isolation, they should continue to influence her when she compromises with others. The compromise strategy I have developed is an intuitive way of extending accepted norms governing single agents to norms governing compromise.

5 Applications

In many situations, it is useful to determine not only what we each individually believe, but what we collectively believe. I have defended a way of understanding the latter. On this understanding, what we collectively believe is what we most value, in a purely epistemic sense. But I have remained neutral on the question of how disagreeing peers should update. Suppose that we are disagreeing peers. Let us grant that we should come to have the same beliefs. It is still a further question whether what you should come to believe—and what I should come to believe—is what we collectively believe, in the sense I have defined.

If disagreeing peers must adopt a single credence, then adopting what they collectively believe seems like a reasonable choice. But even for those who are skeptical about this potential application of the alternative compromise strategy, cases of compromise are not limited to cases in which disagreeing agents trade in their prior credences for matching ones. Several other epistemic situations call for compromising strategies, including some situations involving many agents, and some involving just one.

Even when disagreeing agents retain their individual credences in a proposition, they may still need to act based on a collective opinion. Disagreeing gamblers may need to decide how much of their shared income to bet on a certain horse. Disagreeing weather forecasters may need to report a single opinion to their employers at the radio station.

In another kind of case, disagreeing agents might not willingly trade in their credences, but might still be forced to compromise. For instance, suppose an evil scientist tells us he is going to perform an operation to force us to have the same credence in a certain proposition, and that we can choose only what our new credence will be. I argued earlier that an individual should ask the scientist for whatever alternative credence has the greatest expected epistemic value. Similarly, we should ask for whatever alternative credence maximizes the average of our expected epistemic values. Recall that if we compromise by splitting the difference in our credences, we might end up compromising at one credence when we would both prefer another. It is hard to see how such a compromise strategy could be rationally obligatory.

Strategies for compromise are also relevant to single agents in complicated epistemic situations. For instance, an agent may use a compromise strategy when updating his credence distribution in light of probabilistic evidence, in the sense of Jeffrey 1968. Intuitively, we can get conflicting probabilistic evidence. Suppose a certain screeching bird looks .8 likely to be a bald eagle, but sounds only .6 likely to be a bald eagle. If you had only seen the bird, it would have been clear how to respond: give .8 credence to the proposition that it is a bald eagle. If you had only heard the bird, it would have been rational to give exactly .6 credence to this proposition. But what credence should you assign in the face of conflicting evidence?

In cases where your visual and aural evidence conflict, your scoring rule may determine how you should combine them. Here is a procedure: suppose there are two agents with exactly your credences, except that one agent updates only on your visual evidence, and the other only on your aural evidence. Determine how those agents would perfectly compromise their opinions. Update by accepting their consensus opinion as your own. To take another example, suppose one weather forecaster says it is .6 likely to rain, and an equally trustworthy forecaster says it is .8 likely. In this case, you may update by maximizing the average of the various expected values you would assign to prospective credence distributions, after updating on information from only one of the forecasters.

Finally, some concerns raised in Elga 2007 demonstrate that strategies for compromise may also be relevant to agents with imprecise credences. Elga has a negative conclusion in mind: he aims to demonstrate that if an agent is rational, her credences must be perfectly precise. I am sympathetic with Elga's conclusion, but the arguments needed to establish this claim are more complicated than Elga suggests.

For simplicity, let us suppose that we can represent the belief state of an agent with imprecise credences by a set of probability measures.[15] Elga assumes that an agent with imprecise credences may rationally refuse a bet which is acceptable from the point of view of some but not all probability measures in the set of measures representing her belief state.[16] In other words, if we think of an agent with imprecise credences as if she had a mental committee of agents with precise credences, she may refuse any bet which is unacceptable to any member of her mental committee. But because an agent with imprecise credences may rationally refuse such a large variety of bets, she may rationally refuse a sequence of bets that provides her with an opportunity to win sure money. Foregoing sure money is irrational behavior. So a rational agent must have precise credences.

One could respond to Elga as follows: we should take seriously the suggestion to think of an agent with imprecise credences as if she had a mental committee of agents with precise credences.[17] Such an agent should act in whatever ways a committee should. Elga says that if your mental committee is not unanimously in favor of refusing a bet, "the natural thing to say is that rationality counts the bet as optional for you" (7). But it does not really seem so natural to think that a committee may refuse any bet, as long as just one of its members should prefer to do so. It is more natural to think that in choosing whether to refuse bets, a group should use some compromise strategy to construct a common credence distribution, and then act on the basis of this constructed opinion. Hence we should not accept Elga's premise that a rational agent with imprecise

15. See for instance Tintner 1941, Smith 1961, Levi 1980, Jeffrey 1983, Joyce 2005, van Fraassen 2006.

16. Several advocates of imprecise credences endorse this "conservative" betting strategy. See Williams 1976, Kaplan 1996, Walley 1991, and an extensive catalog of ongoing research at `http://www.sipta.org`.

17. One might also respond that an agent's refusing certain bets constrains what other bets she may rationally refuse. Elga 2007 addresses this response. I will set it aside here.

credences may rationally refuse any bet which is acceptable from the point of view of only some members of her mental committee. By acting on a compromise of her mental committee's opinions, a rational agent could be rationally compelled to accept sequences of bets that let her win money, come what may.

Elga briefly considers the possibility that agents with imprecise credences may still be rationally obliged to accept or refuse any particular bet. He suggests that on such a proposal, "the interval-valued probabilities do little, since they "collapse down" to ordinary point-valued probabilities when it comes to imposing constraints on rational action" (8). I think this response is too quick. For instance, suppose that an agent with imprecise credences updates by "point-wise" conditionalization. In other words, suppose that on receiving some evidence, she updates the set of probability measures representing her belief state by conditionalizing each measure on the evidence received. Furthermore, suppose she has a single credence-eliciting scoring rule, and so acts according to a compromise of the opinions of her mental committee members, all of whom share a credence-eliciting scoring rule. Then she will act according to an arithmetic mean of her committee members' credence distributions.[18] But taking an arithmetic mean of several distributions does not commute with conditionalizing those distributions on a given proposition. So this kind of agent with imprecise credences will act at each stage as if she has a precise credence distribution, but will not act over time as if she has a precise credence distribution that she updates by conditionalization.

This result contradicts Elga's suggestion that if an agent with imprecise credences faces stringent obligations to accept or reject particular bets, her credences may as well be precise "when it comes to imposing constraints on rational action." But the result is friendly towards Elga's overall conclusion. I have argued that as long as an agent has a single scoring rule

18. Here I take an agent's belief state to be represented by a finite set of measures. In this case, the above result follows from the Corollary proved in the appendix. Representing an agent's belief state by an infinite set of measures creates an additional problem: how to parameterize the space of her committee's distributions when calculating their arithmetic mean. This problem resembles Bertrand's paradox. Some theorists treat credences as imprecise chiefly in order to *avoid* paradoxes of this kind. Such theorists have an additional reason to reject the betting strategy currently under consideration.

and updates by "point-wise" conditionalization, she will not act as if she is updating a single precise credence distribution by conditionalization. For instance, she will be subject to diachronic Dutch Books.[19] If that means she is irrational, then we already have a limited version of the conclusion Elga is after, which is a step towards demonstrating the irrationality of imprecise credences.[20]

19. See Teller 1973, Lewis 1999.

20. Thanks to Allan Gibbard, Teddy Seidenfeld, Bob Stalnaker, Roger White, the 2008 MIT-ing of the Minds Conference, and the 2008 Formal Epistemology Workshop for helpful feedback on earlier versions of this paper.

6 Appendix

Example 1. Running the following *Mathematica* notebook verifies that agents with different credence-eliciting scoring rules may not maximize the average of their expected epistemic values in an alternative credence distribution by splitting the difference in their credences.

```
b1[x_] := 1 - (1 - x)^2; b0[x_] := 1 - x^2
l1[x_] := x; l0[x_] := x + Log[1 - x]
ev[x_, m_, f1_, f0_] := (m*f1[x]) + (1 - m)*f0[x]
CDev[x_, m_, f1_, f0_]
  := ev[x, m, f1, f0] + ev[(1 - x), (1 - m), f1, f0]
avgCDev[x_, m_, m1_, m0_, n_, n1_,  n0_]
  := .5 (CDev[x, m, m1, m0] + CDev[x, n, n1, n0])
Maximize[{avgCDev[x, .7, b1, b0, .1, l1, l0], 0 <= x <= 1}, x]
```

For an agent with the Brier score and credence .7 in p, and an agent with scoring rule l_1, l_0 and credence .1 in p, having approximately credence .392965 in p maximizes expected epistemic value:

`{0.924402, {x -> 0.392965}}`.

Example 2. Running the previous notebook with the following additions verifies that agents with non-credence-eliciting scoring rules may not maximize the average of their expected epistemic values in an alternative credence distribution by splitting the difference in their credences.

```
c1[x_] := 1 - (1 - x)^3; c0[x_] := 1 - x^3
Maximize[{avgCDev[x, .7, c1, c0, .1, c1, c0], 0 <= x <= 1}, x]
```

For agents who share the Brier cubed score and have credences .7 and .1 in p, having approximately credence .44949 in p maximizes expected epistemic value:

`{1.75755, {x -> 0.44949}}`.

Example 3. Running the previous notebook with the following addition verifies that the credence distribution preferred by agents sharing a scoring rule in Example 2 is the credence distribution a hypothetical agent would prefer, if she shared that same scoring rule and had the arithmetic mean of their credences:

```
Maximize[{CDev[x, .4, c1, c0], 0 <= x <= 1}, x]
```

For an agent with the Brier cubed score and credence .4 in p, having approximately credence .44949 in p maximizes expected epistemic value:

```
{1.75755, {x -> 0.44949}}.
```

Theorem. If a finite number of agents each use a scoring rule that differs by no more than a constant from a single "shared" scoring rule, then they maximize the weighted average of the expected values they give to a consensus credence distribution by choosing the distribution that a hypothetical agent with their shared scoring rule would prefer, if she were to have that same weighted average of their credences.[21]

Proof. Let us say there are n agents, and that for all $i \in [1, n]$, the ith agent uses scoring rule g_1^i, g_0^i and has credence distribution δ_i over the algebra with atomic propositions p and $\neg p$.

There exist coefficients c_i such that $1 = \sum_{i=1}^{n} c_i$ and such that the following is the weighted average of the expected epistemic values that the compromising agents give to a consensus credence distribution D over the algebra with atomic propositions p and $\neg p$:

$$
\begin{aligned}
&\textsc{waev}(D, \delta_1, g_1^1, g_0^1, \ldots, \delta_n, g_1^n, g_0^n) \\
&= \sum_{i=1}^{n} c_i \textsc{ev}(D, \delta_i, g_1^i, g_0^i) \\
&= \sum_{i=1}^{n} [c_i \textsc{ev}(D(p), \delta_i(p), g_1^i, g_0^i) + c_i \textsc{ev}(D(\neg p), \delta_i(\neg p), g_1^i, g_0^i)] \\
&= \sum_{i=1}^{n} c_i \textsc{ev}(D(p), \delta_i(p), g_1^i, g_0^i) + \sum_{i=1}^{n} c_i \textsc{ev}(D(\neg p), \delta_i(\neg p), g_1^i, g_0^i)
\end{aligned}
$$

21. This result is restricted to credence distributions over four-element algebras, since determining how agents should compromise on such simple credence distributions is sufficient to determine how agents should compromise when they assign different credences to a single proposition. See footnote 5.

By supposition, there is a scoring rule f_1, f_0 such that for any $i \in [1, n]$, there are constants $k_i, l_i \in \mathbb{R}$ such that $g_1^i = f_1 + k_i$ and $g_0^i = f_0 + l_i$. So we can reduce the first summand as follows:

$$\begin{aligned}
&\sum_{i=1}^{n} c_i \textsc{ev}(D(p), \delta_i(p), g_1^i, g_0^i) \\
&= \sum_{i=1}^{n} c_i[\delta_i(p)(f_1(D(p)) + k_i) + (1 - \delta_i(p))(f_0(D(p)) + l_i)] \\
&= \sum_{i=1}^{n} c_i\delta_i(p) f_1(D(p)) + \sum_{i=1}^{n} c_i(1 - \delta_i(p)) f_0(D(p)) + \sum_{i=1}^{n} c_i[\delta_i(p)k_i + (1 - \delta_i(p))l_i] \\
&= \sum_{i=1}^{n} c_i\delta_i(p) f_1(D(p)) + (1 - \sum_{i=1}^{n} c_i\delta_i(p)) f_0(D(p)) + \sum_{i=1}^{n} c_i[\delta_i(p)k_i + (1 - \delta_i(p))l_i]
\end{aligned}$$

This function is simply the sum of the constant $\sum_{i=1}^{n} c_i[\delta_i(p)k_i + (1 - \delta_i(p))l_i]$ and the expected value that an agent with the scoring rule f_1, f_0 gives to credence $D(p)$ in p, when she has credence $\sum_{i=1}^{n} c_i\delta_i(p)$ in p.

Similarly, the second summand, $\sum_{i=1}^{n} c_i \textsc{ev}(D(\neg p), \delta_i(\neg p), g_1^i, g_0^i)$, is the sum of the constant $\sum_{i=1}^{n} c_i[\delta_i(\neg p)k_i + (1 - \delta_i(\neg p))l_i]$ and the expected value that an agent with the scoring rule f_1, f_0 gives to having credence $D(\neg p)$ in $\neg p$, when she has credence $\sum_{i=1}^{n} c_i\delta_i(\neg p)$ in $\neg p$.

The hypothetical agent has credence $\sum_{i=1}^{n} c_i\delta_i(\neg p)$ in $\neg p$ just in case her credence $C(p)$ in p is as follows:

$$\begin{aligned}
C(p) &= 1 - C(\neg p) \\
&= 1 - \sum_{i=1}^{n} c_i\delta_i(\neg p) \\
&= \sum_{i=1}^{n} c_i - \sum_{i=1}^{n} c_i(1 - \delta_i(p)) \\
&= \sum_{i=1}^{n} c_i\delta_i(p).
\end{aligned}$$

So the second summand is the sum of a constant term and the expected value that an agent with the scoring rule f_1, f_0 gives to having credence $D(\neg p)$ in $\neg p$, when she has credence $\sum_{i=1}^{n} c_i\delta_i(p)$ in p.

Hence the initial value $\textsc{waev}(D, \delta_1, g_1^1, g_0^1, \ldots, \delta_n, g_1^n, g_0^n)$ is the sum of a constant term and two expected values: the expected value that an agent with the scoring rule f_1, f_0 gives to credence $D(p)$ in p and the expected value that she gives to credence $D(\neg p)$ in $\neg p$, when she has credence $\sum_{i=1}^{n} c_i \delta_i(p)$ in p. In other words, $\textsc{waev}(D, \delta_1, g_1^1, g_0^1, \ldots, \delta_n, g_1^n, g_0^n)$ is the sum of a constant term, and the expected value that an agent with the scoring rule f_1, f_0 gives to the *credence distribution* D, when she has credence $\sum_{i=1}^{n} c_i \delta_i(p)$ in p.

Since $\textsc{waev}(D, \delta_1, g_1^1, g_0^1, \ldots, \delta_n, g_1^n, g_0^n)$ and the expected value of an agent with credence $\sum_{i=1}^{n} c_i \delta_i(p)$ in p differ only by the addition of a constant term, these functions are maximized at the same values. So agents maximize the weighted average of the expected values they give to a consensus credence distribution, i.e. $\textsc{waev}(D, \delta_1, g_1^1, g_0^1, \ldots, \delta_n, g_1^n, g_0^n)$, by choosing the distribution that a hypothetical agent with their shared scoring rule would prefer, if she were to have that same weighted average of the compromising agents' credences, i.e. $\sum_{i=1}^{n} c_i \delta_i(p)$. □

Corollary. If a finite number of agents share a credence-eliciting scoring rule, then they maximize the weighted average of the expected values they give to their consensus credence distribution by choosing the credence distribution that assigns p that same weighted average of their credences.

Proof. If agents share a scoring rule, then they each use a scoring rule that differs by no more than a constant (namely, 0) from a single "shared" scoring rule. So by the above Theorem, the agents maximize the weighted average of the expected values they give to a consensus credence distribution by choosing the distribution that a hypothetical agent with their shared scoring rule would prefer, if she were to have that same weighted average of the compromising agents' credences.

If the shared scoring rule is credence-eliciting, then the hypothetical agent will prefer her own credences in p and $\neg p$ over any other credences in those propositions. So she will prefer her own credence distribution over any other. Hence agents with a shared credence-eliciting rule maximize the weighted average of the expected values they give to their consensus credence distribution by choosing the credence distribution that assigns p that same weighted average of their credences. □

References

ALSTON, W. 2005. *Beyond "Justification": Dimensions of Epistemic Evaluation*. Cornell University Press, Ithaca.

BRIER, GLENN W. 1950. "Verification of Forecasts Expressed in Terms of Probability." *Monthly Weather Review*, vol. 78 (1): 1–3.

CHRISTENSEN, DAVID. 2007. "Epistemology of Disagreement: the Good News." *Philosophical Review*.

ELGA, ADAM. 2006. "Reflection and Disagreement." Forthcoming in *Nous*.

—. 2007. "Subjective Probabilities Should Be Sharp." Ms., Princeton University.

FITELSON, BRANDON & DAVID JEHLE. 2007. "What is the 'Equal Weight View'?" Ms., University of California, Berkeley. Available at `http://fitelson.org/ew.pdf`.

VAN FRAASSEN, BAS C. 2006. "Vague Expectation Value Loss." *Philosophical Studies*, vol. 127: 483–491.

FRENCH, S. 1985. "Group Consensus Probability Distributions: A Critical Survey." In *Bayesian statistics 2*, J. M. BERNARDO, M. H. DEGROOT, D. V. LINDLEY & A. F. M. SMITH, editors, 183–197. North-Holland, Amsterdam.

GENEST, CHRISTIAN & JAMES ZIDEK. 1986. "Combining Probability Distributions: A Critique and an Annotated Bibliography." *Statistical Science*, vol. 1: 114–135.

GIBBARD, ALLAN. 2006. "Rational Credence and the Value of Truth." In *Oxford Studies in Epistemology*, TAMAR SZABÓ GENDLER & JOHN HAWTHORNE, editors, vol. 2. Oxford University Press, Oxford.

GOOD, I. J. 1952. "Rational Decisions." *Journal of the Royal Statistical Society, Series B*, vol. 14 (1): 107–114.

JEFFREY, RICHARD C. 1968. "Probable Knowledge." In *Probability and the Art of Judgment*, 30–43. Cambridge University Press, Cambridge.

—. 1983. "Bayesianism With a Human Face." In *Testing Scientific Theories*, John Earman, editor. University of Minnesota.

—. 1986. "Probabilism and Induction." *Topoi*, vol. 5: 51–58.

Joyce, James M. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science*, vol. 65 (4): 575–603.

—. 2005. "How Probabilities Reflect Evidence." *Philosophical Perspectives*, vol. 19: 153–178.

—. 2007. "Epistemic Deference: The Case of Chance." *Proceedings of the Aristotelian Society*, vol. 107 (1): 187–206.

Kaplan, Mark. 1996. *Decision Theory as Philosophy*. Cambridge University Press, Cambridge.

Kelly, Tom. 2007. "Peer Disagreement and Higher Order Evidence." Manuscript, Princeton University. Available at http://www.princeton.edu/~tkelly/papers.htm.

Levi, Isaac. 1980. *The Enterprise of Knowledge*. MIT Press, Cambridge.

Lewis, David K. 1999. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge.

Lynch, M. forthcoming. "The Values of Truth and the Truth of Values." In *Epistemic Value*, A. Haddock, A. Millar & D. H. Pritchard, editors. Oxford University Press, Oxford.

Oddie, Graham. 1997. "Conditionalization, Cogency, and Cognitive Value." *British Journal for the Philosophy of Science*, vol. 48 (4): 533–41.

Percival, Philip. 2002. "Epistemic Consequentialism." *Proceedings of the Aristotelian Society*, vol. 76 (1): 121–51.

Savage, Leonard. 1971. "Elicitation of Personal Probabilities and Expectations." *Journal of the American Statistical Association*, vol. 66: 783–801.

Shogenji, Tomoji. 2007. "A Conundrum in Bayesian Epistemology of Disagreement." Ms., Rhode Island College. Available at http://socrates.berkeley.edu/~fitelson/few/few_07/shogenji.pdf.

Smith, Cedric. 1961. "Consistency in Statistical Inference and Decision." *Journal of the Royal Statistical Society Series B*, vol. 23: 1–37.

Teller, Paul. 1973. "Conditionalization and Observation." *Synthese*, vol. 26: 218–258.

Tintner, Gerhard. 1941. "The Theory of Choice Under Subjective Risk and Uncertainty." *Econometrica*, vol. 9: 298–304.

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

Weatherson, Brian. 2007. "Disagreeing about Disagreement." Ms., Department of Philosophy, Cornell University. Available at `http://brian.weatherson.org/DaD.pdf`.

Williams, P. M. 1976. "Indeterminate Probabilities." In *Formal Methods in the Methodology of Empirical Sciences*, M. Przelecki, K. Szaniawski & R. Wojcicki, editors, 229–246. Reidel, Dordrecht.