Episteme, 10, 1 (2013) 87–97 © Cambridge University Press
doi:10.1017/epi.2013.4

EPISTEMIC VALUE AND EPISTEMIC COMPROMISE: A REPLY TO MOSS

AMIR KONIGSBERG
amirkonigsberg@gmail.com

ABSTRACT

Sarah Moss has recently suggested that when they encounter conflict, epistemic peers should not split the difference between the credence that they each assign to some disputed proposition p, as has been suggested by conciliatory approaches to belief revision in the debate surrounding disagreement in the literature. Moss contends that an epistemic compromise between peers need not be the arithmetic mean of prior credences, in the sense that if my credence in some proposition p is x and yours is y, the credence that is the result of our compromise need not be $(x + y)/2$. More generally, Moss's proposal advocates an approach to how estimations of truth value, exhibited in credences, should in fact be considered in resolving conflict and disagreement. The general idea is that splitting the difference between credences may be inadequate, seeing as agents may assign different epistemic values to different credences. While novel and clearly argued, I think that Moss's proposal fails to provide entirely convincing reasons for abandoning the traditional symmetrical approach to epistemic compromise and for adopting the scoring rule model instead. I demonstrate two problems with the model that Moss advocates.

In a recent paper, Sarah Moss has suggested (Moss 2012) that, when they encounter conflict, epistemic peers should not split the difference between the credence that they each assign to some disputed proposition P, as has been suggested by conciliatory approaches to belief revision in the debate surrounding disagreement in the literature (Elga 2007; Christensen 2007; Matheson 2009). Moss contends that an epistemic compromise between epistemic peers need not be the arithmetic mean of prior credences (Moss 2012: 2); if my credence in some proposition P is x and yours is y, the credence that is the result of our compromise need not be $(x + y)/2$, as an equal weighing of our beliefs would suggest. Instead, Moss proposes an alternative strategy for how epistemic peers *ought* to compromise. The basic idea is that splitting the difference (Christensen 2007) between credences may be inadequate, seeing as agents may assign different *epistemic values* to different credences. It is because of this that a compromise that would only consider credences would lack the *epistemic value* that those credences have for the agents involved.

These differences between how different agents evaluate credences supposedly make for *real* differences with regards to what compromising agents *should* do (Moss 2012: 8). It is on these grounds that Moss contends that a more plausible suggestion for epistemic compromise is that disagreeing agents should maximize their *expected epistemic value*. Roughly, they should maximize the epistemic value that alternative credences have for them. This, in a nutshell, is Moss's suggestion for proper epistemic compromise.

While novel and clearly argued, I think that Moss's proposal fails to provide entirely convincing reasons for abandoning the traditional symmetrical approach to epistemic compromise (e.g. Elga 2007; Christensen 2007) and for adopting the scoring rule model instead. In what follows I demonstrate two problems with the model that Moss advocates. The paper is composed of two sections. In the first section I present and discuss Moss's proposal in context. In the second part of the paper I present and expound on two problems on account of which I find Moss's scoring rules solution for epistemic compromise inadequate.

I

Scoring rules

One of the most important factors to consider in assessing normative tactics of belief revision and compromise is the way in which decision-makers quantify uncertainty. In many cases related in the literature, decision-makers quantify uncertainty by expressing their own beliefs using probabilities. In some cases, they don't express their own beliefs but instead defer to the belief of a counterpart or a source of information who they feel is more capable of making such judgments, or has some information that they don't have. In situations such as these, decision-makers use their counterpart's beliefs as a model for what to believe (Joyce 2007). But in both cases – whether they express their beliefs using probabilities or decide instead to defer to someone else's belief – their choice is rooted in personal judgment. And personal judgment, as we all know, is not infallible. Because it is not infallible, various methods have been devised by which decision-makers can improve their assessments. One such method is the scoring rule.

Traditionally (e.g. Brier 1950), scoring rules measure the inaccuracy of a set of probability assessments.[1] They do so by relating an agent or an entity's expected performance to their actual performance through calibration. From an *ex post* point of view, scoring rules can provide a means by which agents can be evaluated with respect to their predictive abilities (Jose 2009: 264). Conceived of in this way, scoring rules enable decision-makers to adjust and improve their predictive assessments based on calibration of past expectations versus the past performances to which those expectations relate. If, for instance, an agent's past performance shows a 0.3 difference between actual performance and expected performance – i.e. as expressed by credences assigned versus performance, then this function will comprise the scoring rule by which present predictions, or levels of confidence, *should* be adjusted. And by adjusting credences, scoring rules can progressively refine margins of error and achieve greater predictive accuracy.

In her novel contribution to the debate surrounding rational responses to disagreement, Moss suggests that epistemic compromise between disagreeing parties can be achieved using scoring rules. But the scoring rules that Moss has in mind are different than traditional scoring rules. Scoring rules as Moss conceives of them are functions that, roughly speaking, measure epistemic value. They are functions that describe the epistemic value, in

1 For recent work on using scoring rules see: Ehm and Gneiting 2012; Gneiting and Raftery 2007; Jose 2009; Bickel 2007. In specific reference to the use of scoring rules to assess credences see: Gibbard, 2008; Percival 2002.

the agent's eyes, of having credence x in some proposition p, in cases where p ends up being true and in cases where p ends up being false. In the sense in which Moss uses them, scoring rules are functions that measure positive as well as negative epistemic value. In making use of scoring rules construed in this way, Moss is following others who adopt this interpretation (e.g. Gibbard 2008, Joyce 1998, 2007, 2009; Oddie 1997; Percival 2002) and seek to construe epistemology decision-theoretically (the value to be maximized by beliefs or credences being truth-related). I will elaborate more on the notion of epistemic value shortly. For the meantime the important point to note is that Moss contends that 'using scoring rules … agents may compromise by coordinating on the credences that they collectively most prefer, given their epistemic values' (Moss 2012: 2). The fundamental idea advocated is that, instead of compromise being conceived of as the arithmetic mean of prior credences, it should be understood as the mean that maximizes expected epistemic value.

Epistemic value

Moss's discussion of epistemic compromise is located in the wider context of epistemic rationality, in which purely epistemic considerations are needed for the desired form of compromise to be rational.[2] The problem that Moss has with conciliatory approaches to compromise is that they do not take into account the fact that agents may value some credences more than others, in what she regards as a 'purely epistemic sense'. These approaches only look at credences, whereas they should also be looking at the epistemic value that different credences may have. And even though Moss departs from these positions, this is not because she believes that they violate the conditions of epistemic rationality by, for instance, considering *non*-epistemic factors in their proposed tactics for epistemic compromise. She departs from them because she believes that they don't consider *all* the epistemic factors that should be considered in epistemic compromise.

If it is to suggest how agents ought to compromise in a purely epistemic sense, the solution that Moss proposes for rational epistemic compromise must also incorporate only epistemic factors. Moss believes that epistemic value is the component that is lacking in conciliatory approaches. And she also believes that its incorporation is what makes her solution to compromise preferable. Perhaps more importantly, Moss contends that epistemic value is a purely epistemic kind of value. And it is that which consequently allows her proposal for epistemic compromise to remain in the domain of epistemic rationality. Following Gibbard (2008), Moss regards epistemic value as the preference that an agent has for a particular credence in relation to one of the contested goals of epistemic rationality, roughly, that of believing true things. The kind of preference she seems to have in mind is that according to which an agent may care more about some types of credences than about others, on purely epistemic grounds.

Caring about credences

If all agents had monotonic scoring rule functions, they would agree that greater value should be placed on more accurate estimates. Thus if there were no divergence in

2 For more on the notion of epistemic rationality see Armstrong 1973; Harman 1995, 1997; Kornblith 1993; Kelly 2002, 2003; Foley 1987, 1993; Feldman 2000; Wedgwood 2002; Velleman 2000; Nozick 1994; David 2001.

https://doi.org/10.1017/epi.2013.4 Published online by Cambridge University Press

attributions of epistemic value between agents in cases where the estimates in question are estimates of the truth value of true propositions, estimates that are closer to 1 would always be better than those that are closer to 0, and a credence such as 0.8 would be more epistemically valuable than credence 0.7, because of its proximity to certainty. So too, in this case, the difference between credences would be valued similarly by different agents: the *difference* between 0.7 and 0.8 would have the same epistemic value as that between 0.2 and 0.3, and as any other pair of credences differing by 0.1. And this would be true for all agents in all situations. But apparently, agents may not share a uniform evaluative scoring rule for credences, even in purely epistemic terms. In fact, believing true things may mean different things for different agents, depending on the credences which they assign.

For instance, we can envision that the function f which describes A's scoring rule of P's possible credences may be expressed as: $(f_1(0.9) - f_1(0.8)) > (f_1(0.6) - f_1(0.5))$, which means that the difference in epistemic value, for A, between 0.9 and 0.8 is greater than – it is preferred to – the difference between 0.6 and 0.5, even though the arithmetic difference between these expressions is the same. A possible reason for this might be that, for A, the former difference in credences relates to estimates of truth that are closer to being certain than the latter; e.g. $[1 - (f_1(0.9) - f_1(0.8))] > [1 - (f_1(0.6) - f_1(0.5))]$. And it is the latter which A prefers. But we might also imagine another agent, B, whose scoring rule of P's credences is expressed as: $(f_1(0.6) - f_1(0.5)) > (f_1(0.9) - f_1(0.8))$, where the arithmetic difference between lower credences is more epistemically valuable for B than the same arithmetic difference between higher credences that have a greater positive proximity to certainty. In this sense for A to assign credence 0.7 to P is different than for B to do so. This means that, in such a case, the epistemic value of 0.7 for A and for B may be different.

It is because the epistemic value of credences may vary between agents that strategies of compromise between agents ought to take this into account. Not to take this variation in preferences into account would not only be to consider insufficient factors – i.e. only considering credences and not considering their epistemic value – but also to consider misleading factors – because credences on their own, on Moss's conception, do not express the epistemic value which they carry.

As different agents have different scoring rules by which they assign values to credences, a simple, purely arithmetic compromise between the credences of different agents will not constitute a perfect compromise of opinion. As the scoring rules which assign value to credences can vary between agents, a merely arithmetic average of the disputed credences may not amount to what we would expect from a perfect epistemic compromise, contra what some conciliatory approaches would suggest.

Hence if compromise means a settlement of differences or mutual concessions, and expressions of epistemic differences between agents are not only credences but also the epistemic values attributed to credences, then a proper epistemic compromise should include both of these factors. Moreover, it is also plausible to assume that a satisfactory mutual concession between disagreeing parties would be that which would be preferred by both parties. And that which would be preferred would seemingly be that which maximized expected epistemic value.

Accordingly, Moss proposes that the 'perfect' epistemic compromise will be the maximization of the average epistemic value of the credence of a proposition for each of the agents in question. It is, she suggests, the coordination of the agents 'on the credences that they collectively most prefer, given their epistemic values' (Moss 2012: 2).

II

Why the scoring rule solution makes sense

Moss seems to take her scoring rule solution to epistemic compromise to be justified for the following reasons:[3]

1. The scoring rules solution is intuitively reasonable.
2. Because it averages combined epistemic values, the scoring rule solution ensures that each party in the dispute is given *something* – what we might regard as a fair share – that they intrinsically care about.
3. The scoring rules solution ensures that neither party in the dispute gets swamped by the epistemic values that another party grants to the same credences.
4. The process of compromising by maximizing expected epistemic value never recommends a credence such that there is some other particular credence on which all agents would *prefer* to compromise.
5. Compromising by maximizing average expected epistemic value is analogous to an intuitively reasonable method of practical compromise in which the average expected practical value of an outcome is maximized.[4]

Having presented Moss's solution and its advocated benefits, I will now proceed to examine how feasible it really is.

The feasibility of the scoring rules proposal

Moss's proposal is novel, intuitively appealing, lucid and coherent. The paper discusses important distinctions that are worthy of attention in their own right. These include credence-eliciting and non-credence-eliciting scoring rules (Moss 2012: 7–9).[5] But setting the merits of the paper aside, my present concern is with the viability of the model that Moss proposes. More specifically, with the contention that scoring rules provide an understanding of 'how an agent may value certain credences over others, *in a purely epistemic sense*' (Moss 2012: 2; my emphasis),[6] as well as with the way in which Moss takes this idea to apply as a preferred solution for epistemic compromise.

To put things concisely, the problem I have with Moss's proposal is twofold. First, it appears to involve what I would like to refer to as a double counting of epistemic value. And secondly, it isn't clear whether the notion of epistemic value which Moss appeals to actually involves the type of value that it would be acceptable and unproblematic to regard as epistemic. And on account of these two concerns, I believe that Moss's proposal fails to provide entirely convincing reasons for abandoning the traditional symmetrical approach to epistemic compromise (Elga 2007; Christensen 2007) and for adopting the scoring rule model instead.

3 I concur with the three factors that Shultz (n.d.) identifies that Moss appeals to.

4 I am grateful to an anonymous reviewer for drawing my attention to points #4 and #5.

5 See also Gibbard 2008.

6 In this Moss is not alone, and my argument applies to others that hold this view. Nonetheless, my focus in this paper is on Moss's account.

https://doi.org/10.1017/epi.2013.4 Published online by Cambridge University Press

That said, I am not sure that symmetrical approaches such as splitting the difference between credences (Christensen 2007) or granting credences assigned by disagreeing agents equal weight (Elga 2007) are the perfect way to compromise epistemically.[7] Nevertheless, neither do I think that Moss's proposal is sufficiently convincing. And this is what I will focus on demonstrating presently. I shall now set forth the two problems I see in the scoring rules approach to epistemic compromise that Moss advocates.

Double-counting

Credences are truth-value estimates (Joyce 1998), and these estimates are subjective estimates. They are an agent's evaluation of the probability that some proposition is true, or an agent's evaluation of the level of confidence that a belief in such a proposition being true deserves. Because of the partially subjective origins of truth-value estimations, the way in which people value credences, insofar as this value is epistemically significant, should seemingly *already* be incorporated into the process according to which credences are assigned.[8]

It seems plausible to assume that an estimation of the truth value of p on the basis of evidence E is at least partly a matter of the degree to which p seems to be *supported* by E. And the notion of 'support' here entails at least some degree of subjective appreciation of the credence that ought to be assigned to p in light of the evidence. And this would be in addition to the assumption that the credence assigned to p is, or can also be, partly directly responsive to E, above and beyond the way E is subjectively understood.

Hence assigning credence is not something that is wholly a function of the available evidence; it is also a function of what that evidence is taken to mean, and the degree of support it is taken to afford the proposition in question. As such, the confidence that the evidence affords the proposition in question is what leads an agent to assign it a particular credence. And the credences chosen already incorporate the agent's preferences regarding the epistemic value of the range of credences for the proposition in question. Otherwise put, the credence chosen as an assessment of p's viability, or of how much confidence p deserves, derives from the function that expresses epistemic value across a range of credences.

If this is indeed the case, then counting the epistemic value of credences separately from the credences themselves appears both superfluous and distortive for reaching epistemic compromise. In this sense a credence that has been assigned to p does not appear to be distinct from the scoring rule according to which that credence is epistemically evaluated. Rather, p's credence seems to be the *result* of the scoring rule by which it is determined. Let me explain.

If these were traditional scoring rules, for instance a Brier score or a Spherical score, and p's credence was 0.7, the scoring rule would not appear to modify the credence assigned to p. In this case, one could switch the scoring rule for p from f_1 to f_2, for instance

7 See my paper (forthcoming) on problems with symmetrical solutions to disagreement problems.

8 Its seems worth mentioning that if your scoring rule is credence-eliciting – which it arguably must be if you are rational – then incorporating your epistemic value into your assignment of credences will not change which credences you assign to propositions, since adjusting your credences to maximize their expected epistemic value will leave your actual credence alone. I discuss this point later in the paper. I am grateful to an anonymous reviewer for this point.

from a Brier score to a Spherical score, and the credences assigned to p would not change, it would remain 0.7.[9] The scores would, in both these cases, calibrate predictions with regard to p, rather than alter these predictions.

But as Moss relates to them, scoring rules do not measure the performance of predictions made under uncertainty. They instead measure the epistemic value of assigning a range of credences to p (see section I, above); as Moss regards them, scoring rules are functions that describe the epistemic value, in the agent's eyes, of having credence X in some proposition p, in cases where p ends up being true and in cases where p ends up being false.

Hence, on this understanding of the role of scoring rules, the credence that an agent assigns to p should be taken as an expression of the scoring rule used for assigning epistemic value to p. If, for instance, the agent happens to be indifferent to small differences between credences (say, the difference between 0.7 and 0.8), as long as both credences are in the right direction – e.g. with proximity to truth (i.e. 1), not falsehood (i.e. 0) – then this should already be apparent in the credence that has been assigned to p – i.e. whether it has been assigned credence 0.7 or 0.8. And in the same way, if one is indifferent to the difference between credence 0.2 and 0.3 when it comes to the likelihood of falsehoods – i.e. one is indifferent to the difference between 0.2 and 0.3 when these are assigned to p, if p ends up not being true, then this should *already* be apparent in the credence that one has previously chosen to assign. And therefore considering epistemic value *after* it has served its purpose is both superfluous and distortive. The role granted to scoring rules on Moss's account suggests that the credence assigned to p by A derives from the scoring rule *f*, which expresses the epistemic value of that credence in relation to p for A.

It is worth mentioning that at a certain point in her demonstration Moss states that 'your scoring rule and your actual credence in a proposition *p* determine the *expected epistemic value* of your having a particular credence in *p*' (Moss 2012: 4), but regrettably, Moss doesn't take this to imply that assigning a particular credence to p may already be an expression of these two factors – your scoring rule for p and the credence you have assigned to p – taken together.

The idea that credences alone do not express the actual epistemic value of propositions because they only refer to estimated *truth* value and not to *epistemic* value is, consequently, problematic. To relate to the conjunction of p's credence *and* the epistemic value of p's credence according to a scoring rule is to weigh what we typically assume has already been weighed in arriving at prior estimations of truth value that are expressed probabilistically in the credence assigned to p.

Furthermore, what I am suggesting here seems also to extend naturally from the conventional way according to which we understand credences and what they represent. There is a probabilistic range, $0 < C < 1$, from within which credences are assigned. As a matter of convention, what a credence measures is the likelihood that a select proposition is true. And the credence that is granted to p out of the probabilistic range represents the level of confidence in p being true. What *significance* a credence can have for an agent, independently of the agent's assessment of p's credence, does not seem relevant. At least not to an assessment of how confident one ought to be in believing p. If a proposition's level of credence is the indicator according to which one can assess how likely it is that p, then p's credence is what one ought to go by. And if one has two diverging

9 I am grateful to an anonymous reviewer for this point.

https://doi.org/10.1017/epi.2013.4 Published online by Cambridge University Press

credences to go by – for instance, the credence that A has assigned to p and the credence that B has assigned to p, then it is in reference to these credences that p's credence ought to be arrived at, if, that is, an estimation of the likelihood of p is what is being questioned. This seems to be the assessment that proper epistemic compromise wants to attain.

In sum, there seem to be two different matters to consider here. The first relates to how likely it is that a particular proposition is true, where this is commonly expressed by credences in the form of probabilities. The second matter relates to what it means, from the point of view of preferences regarding diverse credences, for a particular proposition to have a particular credence. These seem to be two distinct questions, only the first of which seems relevant to an actual assessment of the designated proposition's likelihood.

Epistemic compromise and the agents to which it applies

Moss's discussion of epistemic compromise and her proposed solution focuses on agents that are somewhere between idealized agents and ordinary human believers. The sense in which these agents are idealized is that their scoring rules trade on a purely epistemic currency. Meaning that the way in which they value credences is a function of their (epistemic) preferences with regard to the implications of the proposition being true, as opposed to any other kind of non-epistemic preferences. The sense in which these agents are more human than idealized is the assumed variation, between agents, in the epistemic value attributed to credences. In this sense the fact that 'different agents may value closer estimates differently' (Moss 2012: 3) means that, even though their preferences relate to epistemic factors, they may nonetheless vary between agents.

But if these were ideal agents, conceived of as some kind of epistemic *homo economicus*, or as *Econs* (Thaler and Sunstein 2008), we would likely assume that they would not only condition on evidence in the same way, but that they would also be impartial to whether the difference between varying credences relates to credences that are closer to certainty or further away from it. Credence is, after all, a probabilistic expression of a particular measure of epistemic viability that is determined by where in the probabilistic range $-\ 0 < C < 1$ – a debated proposition is located, in relation to a particular body of evidence. Moreover, what makes for the value of credence is the relation it bears to the actual truth value of the proposition under consideration; and conventionally, the closer it is to certainty the greater value it has. It seems plausible to assume that these idealized, non-human, agents would value credences *purely* epistemically, such that closer estimates would *always* be better because they would represent greater accuracy. And in this sense, epistemic value would, for such idealized agents, perhaps also be a superfluous notion. Once conditionaliztion on evidence is done in the same way for all agents, credences would also be valued in the same way, in accordance with their proximity to certainty. And in this sense such agents will share uniform scoring rules.

This observation about the kind of agents that Moss is referring to is important. It is important because if scoring rules vary between agents, as do the epistemic values that these agents attribute to credences, then there must be grounds for such variation – there must, in this sense, be a reason why an agent values credence in one way rather than another. Furthermore, if an agent's preferences for certain credences over others are to be grounded in 'purely epistemic concerns', then the grounds for such variation must also be epistemic, and it isn't clear that this is something that Moss's account can accommodate.

Epistemic value considerations

Is epistemic value *epistemically* relevant? Moss believes it is. She observes that different agents may value 'closer estimates' differently; 'you may value having 0.9 credence in a truth *much* more than having 0.8 credence, without valuing 0.6 much more than 0.5'. Moss takes such differences in epistemic evaluation to be *epistemically* significant. And it is this assumption that I wish to challenge.

By epistemic value what Moss has in mind is a preferential framework relating to the probabilistic evaluation of truths. Within this framework some probabilistic evaluations are preferred to others. Notably, Moss contends that this preferential framework is epistemic; suggesting that the type of preference for some probabilistic values over others is an *epistemic* preference. Setting aside for the moment *whether* different agents *do* in fact value credences differently, it seems important to ask whether *the fact* that they do so is *epistemically* relevant when considering the appropriate revisionary tactic for epistemic compromise. More specifically, the question is whether *the fact* that agents tend to value credences differently is relevant to an assessment of what *ought* to be the appropriate revisionary response regarding disputed issues between epistemic agents. This, after all, is what epistemic compromise is conventionally supposed to achieve.

In responding, I will start by addressing a practical problem. If different agents can value the same estimates of truth (credences) in different ways, then comprehending these truth estimations will depend on the epistemic values assigned to them by particular agents, set in place by varying scoring rules. And these scoring rules are what will be needed for deciphering the *actual* truth values of the propositions in question. If, for instance, 0.4 means Φ (some epistemic value) to me and Ω (some epistemic value) to you, our *actual* truth estimates will be $\Phi(0.4)$ and $\Omega(0.4)$, respectively. And this means that if we want to understand the credence assigned by each agent we need to abandon the accepted norm and notation of probabilistic evaluations. And as we each grant the arithmetic expression assigned to P a different epistemic value, then epistemically, the same arithmetic expression means one thing for me and another thing for you. And this is problematic, since we no longer have any common way of assessing probability.

But aside from what the epistemic value model entails for our conventional norms for quantifying uncertainty, another matter we must consider in assessing whether epistemic value is relevant for epistemic compromise is whether the epistemic value on which this model is based can unproblematically be regarded as epistemic. I contend that it can't.

In ordinary real-world assessments under uncertainty, the significance that credences carry may vary in accordance with the kind of consequences that they represent. If, for instance, the truth of a proposition that has what we may regard as significant consequences is debated, a greater measure of caution might be taken in assessing its likelihood, perhaps because of the expected cost of being wrong. Alternatively, a more conservative estimation of the proposition's feasibility might be brought into play. And similarly, if the debatable proposition – for instance, a fact or an event – is marginally significant, less caution may perhaps be exercised in assessing its likelihood, because being wrong in such a case may not lead to any substantial outcome. More generally, an agent might prefer credences that are closer to certainty on some occasions because their truth, in relation to a particular proposition, may be practically or emotionally significant.[10] But similarly, on

10 On the prevalence of this phenomenon see Slovic et al. 2005, 2007; Slovic 2010.

https://doi.org/10.1017/epi.2013.4 Published online by Cambridge University Press

other occasions it is the possibility, rather than the probability, of an event with grave consequences that will determine its significance, even if that event's probability is low.

The consequences of a possibility may often determine its epistemic value; the consequences of a particular proposition being true may lead to certain credences being treated otherwise than they would were they considered in relation to a different proposition being true. In the case of a possibility of grave consequences such as a nuclear attack, even low credences may be granted high epistemic value if the consequences are high.

And if, in such a case, we were seeking the consequential significance of particular credences in relation to a particular proposition, then averaging epistemic value to arrive at epistemic compromise might be the right way to go about it. If, as part of the decision-making process competing options were considered in relation to the pending threat, then compromising on the epistemic value of the threat's probability may be feasible. In such a case what we would arrive at would be the significance of credence X for agents A and B in relation to proposition P. But then it would appear that in this case we may have arrived at a compromise, but not an epistemic one.

On any plausible understanding of it, an *epistemic* compromise entails the incorporation of epistemically relevant considerations in accordance with a higher-level truth-eliciting revisionary strategy that doesn't jeopardize what epistemic compromise is all about. Ultimately, epistemic compromise aims to reach the most appropriate conclusion given the competing epistemic assessments, relating to how likely it is that a debated proposition is true, not to what it means for such a proposition to be true. In view of this, if we are going to profitably adopt the claims that Moss defends we should perhaps set aside scoring rules in favour of some more exact tool; in particular, we should adopt rules that assign different accuracy scores to credences in propositions depending on the subject matter of those propositions.

REFERENCES

Armstrong, D. M. 1973. *Belief, Truth and Knowledge*. 1st edn. Cambridge: Cambridge University Press.

Bickel, J. Eric. 2007. 'Some Comparisons among Quadratic, Spherical, and Logarithmic Scoring Rules.' *Decision Analysis*, 4(2): 49–65.

Brier, Glenn W. 1950. 'Verification of Forecasts Expressed in Terms of Probability.' *Monthly Weather Review*, 78(1): 1–3.

Christensen, D. 2007. 'Epistemology of Disagreement: The Good News.' *Philosophical Review*, 116 (2): 187–217.

David, Marian. 2001. 'Truth as the Epistemic Goal.' In M. Steup (ed.), *Knowledge, Truth, and Duty*, pp. 151–69. New York: Oxford University Press.

Ehm, Werner, and Gneiting, Tilmann. 2012. 'Local Proper Scoring Rules of Order Two.' *Annals of Statistics*, 40(1): 609–37.

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

Feldman, Richard. 2000. 'The Ethics of Belief.' *Philosophy and Phenomenological Research*, 60(3): 667–95.

Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.

——. 1993. *Working without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.

Gibbard, Allan. 2008. 'Rational Credence and the Value of Truth.' In: Tamar Szabo Gendler and John Hawthorne (eds), *Oxford Studies in Epistemology*, vol. 2, ch. 6. New York: Oxford University Press.

Gneiting, Tilmann, and Raftery, Adrian E. 2007. 'Strictly Proper Scoring Rules, Prediction, and Estimation.' *Journal of the American Statistical Association*, 102(477): 359–78.
Harman, Gilbert. 1995. 'Rationality.' In E. E. Smith and D. N. Osherson (eds), *Invitation to Cognitive Science*, pp. 175–212. Cambridge, MA: MIT Press.
——. 1997. 'Pragmatism and Reasons for Belief.' In C. B. Kulp (ed.), *Realism/Antirealism and Epistemology*, pp. 123–48. Lanham, MD: Rowman & Littlefield.
Jose, Victor. 2009. 'A Characterization for the Spherical Scoring Rule.' *Theory and Decision*, 66(3): 263–81.
Joyce, James. 1998. 'A Nonpragmatic Vindication of Probabilism.' *Philosophy of Science*, 65(4): 575–603.
——. 2007. 'Epistemic Deference: The Case of Chance.' *Proceedings of the Aristotelian Society*, 107 (2): 187–206.
——. 2009. 'Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.' In Franz Huber and Christoph Schmidt-Petri (eds), *Degrees of Belief*, pp. 263–97. Amsterdam: Springer.
Kelly, Thomas. 2002. 'The Rationality of Belief and Some Other Propositional Attitudes.' *Philosophical Studies*, 110(2): 163–96.
——. 2003. 'Epistemic Rationality as Instrumental Rationality: A Critique.' *Philosophy and Phenomenological Research*, 66(3): 612–40.
Konigsberg, Amir. Forthcoming. 'The Problem with Uniform Soliutions to Peer Disagreement.' *Theoria*.
Kornblith, Hilary. 1993. 'Epistemic Normativity.' *Synthese*, 94(3): 357–76.
Matheson, Jonathan. 2009. 'Conciliatory Views of Disagreement and Higher-Order Evidence.' *Episteme*, 6 (Oct.): 269–79.
Moss, Sarah. 2012. 'Scoring Rules and Epistemic Compromise.' *Mind*, 120(480): 1053–69.
Nozick, Robert. 1994. *The Nature of Rationality*. New edn. Princeton: Princeton University Press.
Oddie, Graham. 1997. 'Conditionalization, Cogency, and Cognitive Value.' *British Journal for the Philosophy of Science*, 48(4): 533–41.
Percival, Philip. 2002. 'Epistemic Consequentialism: Philip Percival.' *Aristotelian Society Supplementary Volume*, 76(1): 121–51.
Schulz, Armin. N.d. 'Epistemic Compromise between Difference Splitting and Scoring Rules: A response to Moss.' Unpublished manuscript.
Slovic, Paul. 2010. *The Feeling of Risk: New Perspectives on Risk Perception*. 1st edn. New York: Routledge.
——. Peters, Ellen, Finucane, Melissa L., and MacGregor, Donald G. 2005. 'Affect, Risk, and Decision Making.' *Health Psychology: Official Journal of the Division of Health Psychology, American Psychological Association*, 24(4 suppl.): S35–40.
——. Finucane, Melissa L., Peters, Ellen, and MacGregor, Donald G. 2007. 'The Affect Heuristic.' *European Journal of Operational Research*, 177(3): 1333–52.
Thaler, Richard H., and Sunstein, Cass R. 2008. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven: Yale University Press.
Velleman, David. 2000. 'On the Aim of Belief.' In David Velleman (ed.), *The Possibility of Practical Reason*, pp. 244–81. Oxford: Oxford University Press.
Wedgwood, Ralph. 2002. 'The Aim of Belief.' *Philosophical Perspectives*, 16(s16): 267–97.

AMIR KONIGSBERG is a member of the Center for the Study of Rationality and Interactive Decision Theory at the Hebrew University in Jerusalem and a Senior Research Scientist at General Motors. His main research interests are in epistemology, decision-making, rationality, and aesthetics.

https://doi.org/10.1017/epi.2013.4 Published online by Cambridge University Press