

## BOOK REVIEWS

Richard Pettigrew, *Accuracy and the Laws of Credence*. Oxford: Oxford University Press (2016), 256 pp., \$74.00 (cloth).

Bayesians hold that rational credences are probabilistic and that rational updating goes by conditionalization. Some add that rational credences defer to known chances and, in cases of ignorance, are spread evenly over possible eventualities. In his ambitious book *Accuracy and the Laws of Credence*, Richard Pettigrew argues that considerations of accuracy are sufficient to underwrite roughly each of these broadly Bayesian principles. The argument strategy employed to make this case is decision theoretic: view credences as options in epistemic decision problems where utility is given by accuracy and then apply appropriate decision theoretic principles to derive the desired Bayesian norms. This style of argument has grown increasingly popular among formal epistemologists in recent years, and Pettigrew's book is without doubt the fullest and most careful exploration of its potential published to date.

The book is neatly divided into four parts, with each dedicated to arguing for a particular constraint on rational credences. While the arguments rely in part on various mathematical theorems, the proofs of these theorems are helpfully placed in appendixes to the main sections of the book, so that formal technicalities rarely interrupt the general flow of Pettigrew's arguments. In its style and organization, the book is a model of clear philosophical writing, striking a smooth balance between mathematical precision and readable prose.

Part 1 (chaps. 1–7) is the longest section of the book and is devoted to arguing for *Probabilism*, the view that rational credences must conform to the probability axioms. In the first chapter, Pettigrew briefly sketches his argument for Probabilism, laying it out in premise-by-premise form. The argument's first premise is *Veritism*, which identifies epistemic utility with accuracy. This premise underlies all of Pettigrew's arguments throughout the book, and its fecundity as a first principle in epistemic utility theory is taken as its main source of support. The second premise posits the *Brier Score* as the correct measure of the inaccuracy of a credence function, according to which the inaccuracy of a credence function is its squared Euclidean distance from the *ideal credence function* that assigns maximal credence to all truths and minimal credence to all falsehoods. The third premise is a dominance principle, which says that if an option is dominated by another option (i.e., yields less utility no matter which epistemically possible world is actual), then the former option is irrational. The final premise in the argument is a

For permission to reuse, please contact [journalpermissions@press.uchicago.edu](mailto:journalpermissions@press.uchicago.edu).

mathematical theorem stating that every nonprobabilistic credence function is accuracy dominated (according to the Brier Score) by a probabilistic credence function and that no probabilistic credence function is so dominated by another credence function. Together, the four premises imply Probabilism.

After defending an appropriate version of the dominance principle (chap. 2), Pettigrew spends the next two chapters of the book considering how to measure the inaccuracy of credence functions. This is where we find support for the second premise of the argument for Probabilism, although the material covered here is also crucial to other arguments throughout the book. In chapter 3, he reviews and argues against some previously proposed constraints on legitimate measures of inaccuracy, while in chapter 4, he axiomatizes his own constraints on inaccuracy measures. As it turns out, the six axioms ultimately defended by Pettigrew jointly imply that every legitimate inaccuracy measure is either the Brier Score or a linear transformation of it.

Pettigrew moves on in chapters 5 and 6 to consider a couple prominent objections to accuracy-based arguments for Probabilism: the Bronfman objection and Howson's robustness objection. His responses here (especially to Bronfman's worry) are thoughtful and neatly worked out. Chapter 7 concludes part 1 with a helpful overview sketching a now fully precise version of the original argument for Probabilism.

Despite the various merits of Pettigrew's defense of Probabilism, we suspect many readers will find the characterization of legitimate inaccuracy measures given in chapter 4 to be the least satisfying aspect of part 1. In particular, we believe more work needs to be done to justify several of the axioms proposed by Pettigrew in his argument for using the Brier Score to measure inaccuracy. Here we highlight one such axiom: *Divergence Additivity*. According to this principle, roughly put, the inaccuracy of a credence function is the sum of the inaccuracies of the particular credences that constitute it. Pettigrew claims that accepting Additivity is a natural thing to do, since a credence function is not a single unified state. Rather, a credence function is a mere collection of individual credences. The members of the collection are unified, but the collection itself is not. Granting Pettigrew's point that a credence function is not a unified doxastic state, it is still not apparent why one should think that Additivity captures the right way to measure the inaccuracy of a credence function. Why not aggregate the inaccuracies of individual credences in some other way, using some other mathematical operation? Additivity is not an implausible assumption to us in this context, but we wish more could have been said in its defense.

In parts 2 and 3, Pettigrew argues for David Lewis's *Principal Principle* and the *Principle of Indifference* (PoI), respectively. Perhaps the more controversial of these arguments will be the one for the PoI. Long a point of contention among Bayesians, the PoI posits that, in the absence of any information about the world, an agent ought to assign equal probability to each of the

worlds she entertains as possible. Hence, given the assumptions of probabilism and a finite set of possible worlds, the PoI entails that an agent's initial or prior credence in any proposition ought to match the proportion of worlds in which the proposition is true. While this fixing of a unique prior credence has appealed to some Bayesians hoping to escape the subjectivism often associated with their school, the justification of the PoI has been a matter of lively debate, to which the argument Pettigrew offers is an original and welcome contribution.

As with the argument for Probabilism, Pettigrew's defense of the PoI rests on the application of a decision theoretic principle to an epistemic decision problem in which the options are credences and utility is accuracy. The decision theoretic principle employed this time is *maximin*, a principle that directs agents to select an option that offers the greatest minimum payoff. (This principle is, of course, familiar to political philosophers as the decision procedure endorsed by John Rawls for use in comparing proposed principles of justice behind the 'veil of ignorance'.) In Pettigrew's epistemic context, *maximin* directs an agent selecting her initial credence function to opt for one that guarantees her the highest minimum accuracy.

Pettigrew proves that the uniform prior (i.e., the credence that assigns equiprobability to each world deemed possible by an agent) is uniquely recommended by *maximin*, when utility is accuracy and accuracy is measured by any plausible scoring rule. While this argument provides a novel and intriguing motivation for the PoI, it cannot be any stronger than the decision theoretic principle on which it rests, and we suspect many readers will find *maximin* a tough sell. To endorse *maximin* is to mandate hyper risk aversion, and one may reasonably wonder why such an attitude should be required of rational agents adopting an initial credence function. Why not rather prefer prior credences according to a risk-loving rule like *maximax* and hence plump for one of the potentially omniscient credence functions? Or why not take a more moderate approach to the whole affair and opt for one of the myriad rules lying somewhere in between *maximin* and *maximax*?

Sensitive to this worry, Pettigrew turns in chapter 13 to explore various other risk-sensitive decision principles one might employ in selecting an initial credence function. The proposals considered here are interesting and crisply explained; however, their philosophical motivation is brief, and they invite further investigation before one could confidently insist on any of them as rationality principles akin to the dominance principles invoked elsewhere in the book.

If Bayesians see Probabilism as the central *synchronic* norm governing an agent's credences at a time, they view *Conditionalization* as the central *diachronic* norm governing an agent's credal changes across time. In the final section of the book, Pettigrew turns to consider whether accuracy arguments

can justify Conditionalization or any other purported norms of diachronic rationality. Pettigrew's answer (defended in chap. 15) is largely negative; he concludes there are no genuinely diachronic norms governing rational credences. Nonetheless, he does argue that a principle closely related to Conditionalization can be bought by accuracy arguments, *Plan Conditionalization* (chap. 14).

According to Plan Conditionalization, the only rational way to plan to update in the face of new evidence is by conditioning on the evidence received in line with the familiar ratio rule.<sup>1</sup> Pettigrew supplies three arguments for this principle. The principal of these arguments, stemming from earlier work by Greaves and Wallace, takes possible updating plans (i.e., functions from a partition of potentially learnable propositions to credence functions), rather than single credence functions, as options to evaluate. (The accuracy of an update rule at a given world then being defined as the accuracy of the credence function it recommends at that world.) The decision-theoretic principle deployed in the argument is simply *expected utility maximization*: an agent should prefer an updating plan that minimizes expected inaccuracy by the lights of her current credence function. Granting this principle, since it can be shown that conditionalization minimizes expected inaccuracy among possible updating plans relative to an agent's given probabilistic credence function, Plan Conditionalization can be established.

Readers who, like us, find Pettigrew's skepticism concerning diachronic rationality questionable will find it interesting that Pettigrew provides a theorem that comes close to vindicating full Diachronic Conditionalization. The theorem in question states, very roughly, that any credence function  $c$  will view  $c(\cdot|E)$  as having the least expected inaccuracy among credence functions that assign maximal credence to  $E$  (if the expectation is computed only summing over worlds where  $E$  is true). Pettigrew rejects an inference from this theorem to Diachronic Conditionalization on the grounds that once one has learned new information one's prior credence function can be recognized as defective (it failed to assign probability 1 to the new information, after all) and hence lacks any right to offer a privileged recommendation regarding posterior credence. Not all readers need find this concern compelling, however. Those who do not may actually come closer to deriving a fully Bayesian epistemology from considerations of accuracy alone than even Pettigrew does.

The book's final chapter (chap. 16) is one of its many highlights. Pettigrew concludes by surveying the various research questions that remain open for the epistemic utility project. The rich array of open problems iden-

1. Unless, of course, the evidence received has zero prior probability, in which case the ratio rule gives no direction; Conditionalization is compatible with any response to such evidence.

tified by Pettigrew provides not only a fitting conclusion to the book but also excellent motivation to work through it in the first place. As his survey demonstrates, the epistemic utility program intersects with many of the most popular topics in current formal epistemology and philosophical decision theory, including risk sensitivity, infinitesimal probabilities, credences in self-locating propositions, and imprecise credal states.

At the book's conclusion, we do believe there remain serious concerns regarding Pettigrew's project. For example, to the extent that we can grasp the notion of epistemic utility, Veritism strikes us as too counterintuitive a claim to be justified solely by reference to its theoretical fecundity in the context of Pettigrew's arguments. For example, how can it be right to say that an agent with high confidence in two true but trivial propositions and low confidence in one equally true but far deeper proposition is epistemically better off (on any significant way of spelling out that notion) than one with high confidence in the deep proposition and low confidence in the trivial ones?<sup>2</sup> Additionally, Pettigrew's project seems committed to the claim that the rationality of an individual credence is determined by whether it is a constituent of a credence function that has highest epistemic utility for the agent. But just as the epistemic utility of one's future credences, or of another person's credences, is plausibly irrelevant to the rationality of one's current credences, so it also seems plausible that the epistemic utility of one's credences as a whole should be irrelevant to the rationality of one's credence in a particular proposition.

Nonetheless, the progress Pettigrew has made thus far in elegantly constructing and creatively defending his accuracy-only approach to epistemic utility theory more than warrants further consideration of these worries and of the prospects for devising an adequate response to them. In sum, *Accuracy and the Laws of Credence* is the cutting edge in the burgeoning field of epistemic utility theory and is required reading for anyone interested in applications of decision theoretic reasoning to epistemology.

CHAD MARXEN, BROWN UNIVERSITY  
GERARD ROTHFUS, UNIVERSITY OF CALIFORNIA, IRVINE

2. In personal correspondence, Pettigrew has pointed out that many of the book's formal results could still be won without treating all propositions as equal in epistemic importance. The implications of this for epistemic utility theory are well worth exploring, although moving to such an approach may signal something of a retreat from unvarnished Veritism.