

Bayesianism, *Quo Vadis?*—Critical
Notice: David Corfield and Jon
Williamson (eds.), *Foundations of
Bayesianism**

Mathias Risse†

David Corfield and Jon Williamson (eds.), *Foundations of Bayesianism*.
Dordrecht: Kluwer Academic Publishers (2001), 428 pp. \$110.00 (cloth).

This is a review essay about David Corfield and Jon Williamson's anthology *Foundations of Bayesianism*. Taken together, the fifteen essays assembled in the book assess the state of the art in Bayesianism. Such an assessment is timely, because decision theory and formal epistemology have become disciplines that are no longer taught on a routine basis in good philosophy departments. Thus we need to ask: Quo vadis, Bayesianism? The subjects of the articles include Bayesian group decision theory, approaches to the concept of probability, Bayesian approaches in the philosophy of mathematics, reflections on the relationship between causation and probability, the Independence axiom, and a range of criticisms of Bayesianism, among other subjects. While critical of some of the arguments presented in the articles, this review recommends Corfield and Williamson's volume to anyone who is trying to stay abreast of Bayesian research.

Bayesian philosophers disagree about much, but they all endorse the centrality of probabilistic thinking in various areas of inquiry, especially in the philosophy of science and decision theory. The most distinct subgroup are the subjectivists, who conceive of rationality in terms of maximizing expectations composed of subjective probabilities and utilities. Developing these sketchy ideas into theories involves questions encompassing value theory, epistemology, and metaphysics. For this reason, and because

*Received September 2002.

†To contact the author, please write to Mathias Risse, John F. Kennedy School of Government, Harvard University, 79 JFK St. / Taubman 156, Cambridge, MA 02138; e-mail: mathias_risse@ksg.harvard.edu.

Philosophy of Science, 70 (January 2003) pp. 225–231. 0031-8248/2003/7001-0014\$10.00
Copyright 2003 by the Philosophy of Science Association. All rights reserved.

it exemplifies the project of bringing formal reasoning to bear on ever more questions traditionally treated (or non-existent) without it, Bayesianism was involved in debates central to twentieth-century philosophy: debates about the ontology of decision-making, belief-revision (including confirmation and induction), the nature of explanation, scientific progress, the nature of belief and knowledge, rationality, and practical reasoning. In addition, Bayesians engaged in internal debates about the concept of probability, the relationship between probability theory and logic, between probabilistic and causal reasoning (Newcomb's problem), and between decision theory and game theory. Yet while Bayesianism has recently become prominent in statistics, the social sciences, and computer science, silence has fallen over the old battlegrounds on which philosophers once hoisted Bayesian flags. Many philosophy departments ceased to offer classes in decision theory, and in the philosophy of science debates in the philosophy of physics and biology claimed center stage. *Quo vadis, Bayesianism?*

David Corfield and Jon Williamson's volume demonstrates that Bayesianism has lost none of its philosophical interest. Most of these 15 articles were presented at a conference in London in 2000. The spirits of Popper and Lakatos are as present as that of the Reverend Bayes, and thus probably more than some readers welcome. The editors are to be congratulated on getting the collection published quickly. (Yet this speed seems to have come at some cost. The index is put together carelessly: e.g., Galavotti's article fails to contribute any keywords. There are also more typos than Kluwer should appreciate.) The volume offers a panorama of current research in Bayesianism. One lesson is that the reason for Bayesianism's neglect in philosophy cannot be that the interesting questions have been settled. They have not. This neglect is a sociological phenomenon: philosophers happen to be interested in other issues now. Possibly, the interdisciplinary spirit that fueled interest in Bayesianism receded. Possibly, young philosophers are less interested in acquiring additional, seemingly "non-philosophical" qualifications, like working knowledge of probability theory. Be that as it may, let us explore what we learn from these authors about the state of the art in Bayesianism. Apologies to Williamson, Peter Williams, Jeff Paris and Alena Vencovska, and James Cussens, whose contributions are neglected due to this reviewer's relative ignorance.

First, a complaint: Philippe Mongin's article on the "Paradox of the Bayesian Experts" was not presented at the conference, but the editors wisely chose to include it. But it might have been better to include Mongin's seminal 1995 article on Bayesian aggregation, rather than this 1998 follow-up discussing the state-dependent case. The editors introduce this piece as one of those that "scrutinize the acceptability of particular axioms" (11), in this case Pareto. Mongin does show that Pareto is inconsis-

tent with plausible ideas about aggregating utilities and probabilities in Bayesian group settings. However, his article does much more. Mongin and others inaugurated a new research program: Bayesian group decision theory. It is unfortunate that the editors do not consider this area at greater length. Bayesian group decision theory might add to the renewal of interest in Bayesianism at least among those philosophers who think that *individual* decision theory, not Bayesianism as such, has become baroque.

This volume also debates long-standing issues: Colin Howson and Maria Carla Galavotti explore approaches to the concept of probability, Richard Bradley discusses Ramsey's representation theorem, and Edward McClennen reconsiders the Independence axiom. Howson disconnects the idea of probability as partial belief from prudential decision making and thus from utility theory. He suggests that "one should not need a general theory of rational preference in order to talk sensibly about estimates of uncertainty and the laws these should obey" (137). Following Leibniz and James Bernoulli, Howson strives to give the idea of consistent assignments of fair betting quotients a logical meaning, by generalizing the idea of consistency in terms of extendibility of valuations to a complete set of propositions in accordance with the laws of logic. As Howson stresses, this approach runs contrary to the development of probability as partial belief since the seventeenth century. He claims that Bayesian orthodoxy mistakenly developed that notion in terms of "fairness of betting quotients as a lack of differential impact," rather than "fairness of betting quotients as absence of bias," which is the idea he wants to develop as part of logic.

Howson's claim that a theory of uncertainty should not require a theory of rational preferences seems plausible if we disregard the fact that this is uncertainty attached to *beliefs*, rather than to *propositions*. Logic clarifies the relationship between propositions, and generalizing that idea leads to the early Carnap: one way in which propositions relate is by having partially overlapping contents. However, if one conceives of probability in terms of *partial beliefs* of a rational agent, then it is unwise to ignore the connection to rational preference. Denying such a connection is not required by the position that intellectual judgments cannot be *reduced* to value judgments, which is a point Howson champions and that most workers in this field do not reject. Instead, denying this connection amounts to rejecting the idea that rational agency requires even a minimal harmony between beliefs and value judgments of rational agents, and thus to denying that beliefs are action-guiding in any way, which seems misguided. Such a minimal harmony is precisely what Dutch book arguments establish: if one's beliefs conflict with the probability axioms, one can be enticed to act in ways that undermine what one values.

What about Howson's appeal to the distinction between fairness as

absence of bias and fairness as absence of differential impact? There seems nothing wrong with the second idea of fairness he denounces. At any rate, fairness is too amorphous a notion to decide which approach to probability to choose. Finally, Howson's comparison between logic and probability theory is dubious. He proves a theorem stating that an assignment of betting quotients is "consistent" precisely if it satisfies the probability axioms. He submits that this theorem plays the role of a soundness and completeness theorem in logic. Howson captures consistency by reference to the behavior of bookies. Once this is spelled out, it becomes obvious that consistency of betting quotients amounts to satisfying the probability axioms. Not much remains to be proven, and one is left with the question of why exactly partial beliefs should be characterized the way Howson, and the bookies, say they should. Orthodoxy answers this question in terms of Dutch-book arguments. Since Howson dismisses those, he is simply left with the question. Bayesian orthodoxy prevails.

Galavotti draws an attractive picture of the arch-subjectivist de Finetti. She emphasizes that, through his notion of exchangeability and his inquiries into scoring rules, de Finetti can account for the seemingly objective uses of probability. But although his interest in such uses became stronger late in his life, he kept insisting that frequencies, like symmetry considerations, are limited to being mere ingredients of probability evaluations. Yet once they are granted as much, is it such a big step to pluralism about probability and then to assessing the relationship of the different notions, as done, for example, by Lewis' Principal Principle? Galavotti makes de Finetti look less dogmatic than many thought he was, but his exclusive subjectivism seems to undermine itself precisely by being accommodating. Merely reading Galavotti's piece, one would surmise that both subjective and objective probability have come to stay. Yet de Finetti's subjectivism and, in its footsteps, Richard Jeffrey's radical probabilism have resources to answer this challenge, and in light of those one starts wondering about the state of the art in the *objectivist* camp: is there no representative of "objective chance" insisting that more is left of it than such accommodation into subjectivism?

Bradley reconstructs Ramsey's representation theorem. A difficult author, Ramsey left brilliant but sketchy ideas in need of elaboration. Bradley develops his theorem without glancing over remaining questions, in particular ontological ones. However, I wish to discuss questions that arise about representation theorems in general. Recall why there was interest in representation theorems initially. Von Neumann and Morgenstern, and Savage, attempted to legitimize probability and utility by deriving them from the seemingly more tractable notion of "preferences." Preferences were considered observable and thus, in the heyday of logical empiricism and behaviorism, bestowed the required legitimacy. Nowadays, probab-

ity and utility are not as much in need of justification. The strong connection between scientific usefulness and observability turned out to be hard to defend, and belief and utility are useful notions even if they remain closed to measurement. Furthermore, deriving them from preferences provides a dubious foundation for these notions: observations often fail to license the desired inferences, and preference theories work with broader classes of preferences than we can ever observe. What then is the purpose of representation theorems, now that the original motivations are missing?

This question arises forcefully for Ramsey's theorem. For Ramsey shares a concern with von Neumann and Morgenstern, and with Savage, without worrying about observability. That concern is to make probability and belief *measurable*. Yet Ramsey's theorem rests on "conditional prospects" and "ethically neutral propositions" that resist integration into a theory championing observability. But from what point of view is a *measurement* approach to representation theorems useful if based on *introspection*, as Ramsey's seems to be? (Understanding Ramsey's theorem as useful for third-person inquiries means taking seriously what Jeffrey once called "a bizarre caricature of a psychological test;" cf. Chapter 4 of his *Logic of Decision*). Bradley claims that a theory must make sure its variables can be measured. But we need a bigger theoretical picture. Why must we guarantee such measurability if we do so in a way that is closed to third-person inquiries? Jeffrey's work contains a different take on representation theorems, namely, to think of them as establishing (non-reductive) logical connections among the entities of the theory (i.e., a "logic of decision"). Yet if that motivates Bradley's investigation, one wonders about the insistence on measurement issues. A different response is to integrate representation theorems and decision theory better into the philosophy of practical reasoning. Measurements based on conditional prospects and ethically neutral propositions are conducted in and by the mind, geared towards guiding deliberations, rather than towards making values and beliefs accessible to observers. Such an approach is exciting in light of the view held by philosophers like Jean Hampton and David Wiggins that decision theory is useless to the philosophy of practical reasoning. At any rate, I am far from suggesting that such questions cannot be answered, but more reflection about the status of representation theorems is desirable now that decision theory has matured. However, the fact that Bradley fails to offer more on this score does not distract from the fine job he does in revitalizing interest in Ramsey.

McClennen revisits the Independence axiom, and does so in an illuminating way. That by itself is no small feat given how much ink has been spilled over this postulate. He insists that traditional counterexamples to Independence (Kahneman and Tversky, Allais, Ellsberg) display a phenomenon of complementarity typical of disjunctions. That is, they

show that it can be rational to be influenced in one's attitude towards one option by what else is or could have been available. However, McClennen omits a promising response championed (say) in Jim Joyce's *Foundations of Causal Decision Theory* that still holds against McClennen's sophisticated version of the argument: underdescribed outcomes leave traditional decision theory inadequate. Thus Independence only applies if the outcomes cover everything that matters. The cost of this move is that one needs a notion of outcome detached from its colloquial understanding. However, that is a bullet that any brand of consequentialism can bite. Consequentialists in ethics are keenly aware of that, and at least in this case what is no problem in consequentialist ethics should not be one in decision theory.

The editors divide their articles into four sections ("Bayesianism, Causality, and Networks," "Logic, Mathematics, and Bayesianism," "Bayesianism and Decision Theory," "Criticisms of Bayesianism"). The opening papers display more unity than the other sections, being all devoted to the connections between probability and causation. Judea Pearl contributes the lead article, advocating the approach in his 2000 book *Causality*. Pearl thinks of Bayesianism as an attempt to capture what we know, and argues that it remains incomplete on that score. The probabilistic language must be (non-reductively) enriched by causal vocabulary to come "closer to where knowledge resides" (19). Mathematizing causation, Pearl proposes a language of directed graphs for this purpose. While philosophically informed, Pearl is based in AI and statistics where skepticism of causal vocabulary abounds. For most philosophers nowadays the challenge is not primarily to give legitimacy to causal vocabulary, but to explore how work on causation and work on probabilistic reasoning can be reconciled. Philosophers of causation, especially the highly visible group endorsing a counterfactual analysis, should take better account of Bayesians approaches. Counterfactual reasoning also triggers a debate between Pearl and Philip Dawid. Appealing to Popper, Dawid criticizes Pearl for integrating such reasoning, which he finds "metaphysical" rather than scientific, worrying that it licenses unwarranted inferences. Yet Dawid's uncritical appeal to Popperian methodology is misplaced. Much work has been done to refute it. Even more work has been done in violation of that methodology. Dawid should engage with post-Popperian philosophy before dismissing ideas for being part of it. Does Lewis-style analysis not deserve more than automatic rejection in virtue of not having citizenship in Popper's universe?

We change topics again. David Corfield's piece on "Bayesianism in Mathematics" inaugurates a new research area. There are Bayesian reconstructions of the history of mathematics just as there are Bayesian reconstructions of the history of science. But Corfield also suggests an-

other way of giving Bayesianism a role in mathematics, by providing a notion of inference generalizing deductive inference. While this idea is alien to a traditional understanding of mathematical proof, Corfield succeeds in linking it to mathematical practice. He submits that, now that computer programs play a large role in mathematics, less than obviously deductive proofs have become prevalent. Corfield provides three examples where Bayesianism offers insights into mathematics: by investigating proceeding by analogy, by enumerative induction, or by strategy choice. However, as for analogy and strategy choice, it remains unclear how precisely they relate to Bayesianism. Also, such examples operate at the level of historical reconstruction and fail to investigate whether there is a more general notion of proof than the one treasured by logicians. Nevertheless, Corfield shows that there are rewarding questions here.

Let us proceed to the section on criticisms of Bayesianism. Max Albert presents a theorem showing that under common circumstances a Bayesian observer cannot rule out any behavior as irrational. While this result questions Bayesianism as an explanatory tool, from the standpoint of decision theory advising the first-person perspective, it is hard to interpret. When making a decision, an agent does so in light of her beliefs. Why should she be discouraged from following Bayesian advice if for each available action there exist prior beliefs (most likely not hers) rendering it rational? The remaining two articles revisit the debate between Neyman/Pearson hypothesis testing and Bayesian statistics. Donald Gillies argues that Bayesianism is ill-suited for learning about the framework used in an experiment. Suppose we assume the framework is a Poisson model with parameter λ . Then Bayesianism captures learning about λ , but cannot accommodate learning about the framework itself. One may try to switch to a more abstract Bayesian model, but Gillies rebuts such attempts. The ball is in the Bayesian court, and Gillies' seems a tough challenge to meet. Deborah Mayo and Michael Kruse discuss the accommodation of stopping rules (*viz.*, plans for when to stop an experiment). For Bayesians, unlike Neyman/Pearson advocates, the fact that somebody tried repeatedly to bring about a result makes no difference to the data's evidential input: it is just as relevant as if the agent had planned all along to stop at a certain point. Mayo and Kruse are aware that the traditional Bayesian response is to bite the bullet. They illuminate the costs of this move, but at the end of the day, that simply is what Bayesians can and must do.

Corfield and Williamson's impressive volume is recommended reading for anybody trying to stay abreast of Bayesian research. Bayesianism has moved with full force into the twenty-first century. Hopefully, philosophy departments will not be left behind.