BOOK REVIEW

Bayesian Epistemology by Luc Bovens and Stephan Hartmann, Oxford University Press, 2004, pp. ix + 159. H/b £40.00, ISBN 0-19-926975-0. P/b £17.99, ISBN 0-19-927040-6.

Luc Bovens and Stephan Hartmann (hereafter B&H) have produced an impressive book on probabilistic epistemology. There have been numerous applications of the probability calculus to issues in philosophy of science, but B&H's main interest in this book is problems of traditional epistemology. They employ concepts of probability theory to give precise formulations to some long-standing issues of epistemology and offer their resolutions based on formal derivations. This is a welcome addition to the epistemological literature, for despite the recognition that most empirical evidence supports our beliefs only probabilistically, traditional epistemologists have been slow to take advantage of formal methods of the probability calculus. Chapter 1 of the book ("Information") analyzes the effect of the coherence of evidence on the probability of the propositions it supports; Chapter 2 ("Coherence") proposes a formal measure of coherence; Chapter 3 ("Reliability") analyzes the effect of the coherence of evidence on the probability that the sources are reliable; Chapter 4 ("Confirmation") introduces an additional layer of "testable consequences" between the hypothesis and the evidence to examine the roles of the variety of evidence and the auxiliary hypotheses; Chapter 5 ("Testimony") analyzes the complex way the specificity of concurring reports affects their probability. Anyone with serious interest in formal epistemology will benefit greatly from reading this book not just because of the results obtained, but also because of the concepts and techniques B&H develop for their analyses that have potential applications beyond the issues directly addressed in the book.

Since it requires a serious commitment on the part of the uninitiated reader to decipher formal notations and navigate through the sea of equations, I want to offer a brief description of what issues of traditional epistemology B&H's analyses illuminate most, so as to help lay readers to decide the level of their intellectual investment. The most important theme of the book is the impact of coherence (or more generally the impact of the relation) among independently produced pieces of evidence on the confirmation of the hypothesis. Already at the end of Chapter 1, B&H mention some implications of their formal results for the coherence theory of justification, but there are of course "many versions of the coherence theory of justification" (26) as B&H note, and it is clear that the version B&H have in mind is not the coherence theory of justification as an alternative to foundationalism in traditional epistemology. This is because B&H's model there assumes that a single piece of evidence is partially reliable on its own in the sense that it raises the probability of the proposition it supports at least to some extent. The role of coherence envisaged here is therefore an *enhancement* of extant justification, and not a *creation* of justification from scratch. In later chapters B&H incorporate the

possibility that the evidence may have no relevance to the proposition it (purports to) support, but those models still assign a positive probability to the hypothesis that the source is reliable. As a result, a single piece of evidence provides *some* justification for the proposition it supports. The role of coherence there is still an enhancement of extant justification, and not a creation of justification.

In fact B&H explicitly denies the latter role to coherence: "Clearly, the reports of fully unreliable witnesses should be of no consequence to our degree of confidence regarding the matters attested to, regardless of the coherence of the reports" (14). Those epistemologists who hold or hope that coherence is an independent source of justification—and thus can help solve the classic problem of regress in justification—may be disappointed by this pronouncement and dismiss the book as irrelevant to their interest. That would be unfortunate since complex models of later chapters contain tools that are useful for analyzing more ambitious forms of coherentism than B&H consider in the book (see my "Justification by coherence from scratch" forthcoming in *Philosophical Studies* for one such analysis). There is, meanwhile, no indication that B&H's methods can help us deal with deceiver scenarios familiar to epistemologists, such as the Cartesian demon world and the brain-in-a-vat story. This is because their models assume prior knowledge of independence of the sources of evidence, while in these deceiver scenarios a single source (the deceiver) generates all evidence.

B&H's analyses are valuable in most direct ways for those epistemologists, such as C. I. Lewis, who reject coherentism as an alternative to foundationalism, but who acknowledge its positive role in epistemic evaluation. The reader needs to be aware though that despite the title of its first section ("C. I. Lewis's Heritage"), Chapter 1 does not have much to do with Lewis's own analysis of coherence ("congruence" in his terminology). Lewis's main idea on the subject is that when independently produced reports (pieces of evidence) turn out to be coherent, we can be more confident that the witnesses (the sources of evidence) are reliable; whereas B&H's model in Chapter 1 assumes that our assessment of the reliability of the sources is not affected by the coherence of evidence they produce. It is later in Chapter 3 that B&H drop this assumption and analyze the impact of coherence on our assessment of the reliability of the sources. Also crucial for a proper analysis of Lewis's approach is Chapter 5. As B&H mentions there, an important component of Lewis's view is that unreliable witnesses are unlikely to produce coherent stories because each story "is just one out of so many false stories that could have been told" (112). In other words, the stories are too specific to be coherent by coincidence. The impact of the specificity of concurring reports on their probability is the subject of Chapter 5. In short epistemologists interested in Lewis's view of coherence must read the entire book.

For the remainder of this review I want to focus on one of the most elegant and striking results of the book, the equation (1.5) on page 18, which states (to put it in English) that other than the degree of reliability of the sources, the probability of the conjunction of the propositions upon receiving the evidence is solely dependent on the "weight vector" of the propositions, where the weight vector (defined formally on page 17) is intuitively the extents of the "overlaps" among the propositions. To see what this means by a simple example, think of two independently produced reports, $REPR_1$ and $REPR_2$, whose respective contents are R_1 and R_2. The claim of (1.5) is that other than the degree of reliability of the sources, the probability of the conjunction R_1 & R_2 upon

receiving the reports $REPR_1$ and $REPR_2$—or $P(R_1, R_2|REPR_1, REPR_2)$ in formal notation—is solely dependent on the following three values; $a_0 = P(R_1, R_2)$, $a_1 = P(R_1, \neg R_2) + P(\neg R_1, R_2)$, and $a_2 = P(\neg R_1, \neg R_2)$. It is easy to show that in this particular case these three values are determined completely by the prior probability of the conjunction $P(R_1, R_2)$ and the prior probability of the disjunction $P(R_1 \vee R_2)$. No other part of the probability distribution, such as $P(R_1)$ or $P(R_2)$, is relevant to the probability of the conjunction upon receiving the evidence, so long as $P(R_1, R_2)$ and $P(R_1 \vee R_2)$ remain the same.

B&H combine this striking result with a principle of Bayesian coherentism, BC_1 (stated fully on page 11), which says essentially that coherence is truth-conducive. What follows from these two points is that the degree of coherence is a function of the weight vector. The reasoning is that if the probability of the conjunction upon receiving the evidence is completely determined by the reliability of the sources and the weight vector of the propositions, then the degree of coherence should be completely determined by the weight vector because according to BC_1, coherence is whatever relation is truth conducive, given the same reliability. To apply this result to the simple case of two reports above, the degree of coherence is solely determined by $P(R_1, R_2)$ and $P(R_1 \vee R_2)$, and not affected by any other part of the probability distribution, such as $P(R_1)$ or $P(R_2)$. This is a surprising result in itself, but B&H build on this finding to argue further in Section 1.4 that there is no measure of coherence (in the sense of total ordering) that is truth conducive regardless of the degree of reliability of the sources. The strong form of Bayesian coherentism is therefore untenable.

This is a remarkable result supported by an ingenious reasoning, but there is one technical detail that makes me pause. The point of the equation (1.5) is that *other than the degree of reliability of the sources*, the probability of the conjunction of the propositions upon receiving the evidence is solely dependent on the weight vector of the propositions. Obviously, in order to make this claim B&H must translate the intuitive notion of reliability into the formal language of probability, and my concern is that theirs is not the only possible translation. According to B&H, the degree of reliability, r, is a function of the true positive rate p_i and the false positive rate q_i of the evidence, where the true positive rate is the probability that the evidence is in favor of the proposition given that the proposition is indeed true, or $P(REPR_i|R_i)$, while the false positive rate is the probability that the evidence is in favor of the proposition given that the proposition is actually false, or $P(REPR_i|\neg R_i)$. More specifically, B&H's measure of reliability is $r = 1 - q_i/p_i$. In other words, they measure reliability by subtracting the ratio of the false positive rate to the true positive rate from one, which is quite circuitous.

The adoption of r as their measure of reliability enables B&H to obtain the elegant and striking result (1.5), but the question is whether there is a compelling reason for adopting r. There are many sensible alternatives, including various Bayesian measures of confirmation found in the literature (see Branden Fitelson, "The Plurality of Bayesian Measures of Confirmation and The Problem of Measure Sensitivity," *Philosophy of Science* 66 (1999), pp. S362-S378 for a survey). If we are not too particular about the intuitive sense of "reliability", we may even opt for $r^+ = P(R_i|REPR_i)$, which is simply the probability that the proposition is true given that the evidence is in favor of it. In footnote 10 on page 17, B&H state that their results in Chapters 1 and 2 do not depend on their choice of a reliability measure and promise to show this later, but what they

show later on pages 74-75 is only that the earlier results hold even if the coherence of evidence affects the assessment of reliability, while reliability itself is still measured in the old way from the true positive rate and the false positive rate. Adoption of an alternative measure of reliability would not affect B&H's main points if r were a function of the alternative measure since one could then rewrite (1.5) using the alternative measure and the weight vector, thereby making the probability of the conjunction upon receiving the evidence still a function of the reliability and the weight vector. However, that is not the case with all alternative measures. For example, as the equation below (obtained from (1.5) with n = 1) shows, r is not a function of $r^+ = P(R_i|REPR_i)$ but a function of r^+ *and* the prior probability of the proposition, $P(R_i)$.

$$r = 1 - \frac{P(R_i) \times (1 - r^+)}{(1 - P(R_i)) \times r^+}.$$

This means that if we measure the reliability of the sources by r^+, then the probability of the conjunction upon receiving the evidence depends on the degree of reliability r^+, the weight vector, *and* the prior probabilities of the individual propositions. In the simple case of two reports above, for example, one can no longer say that $P(R_1)$ and $P(R_2)$ are irrelevant to the probability of the conjunction upon receiving the evidence. Since the irrelevance of these prior probabilities is the basis of their refutation of strong Bayesian coherentism, it is crucial that B&H explain why one cannot adopt r^+ or other measure of reliability of which r is not a function so as to keep strong Bayesian coherentism alive. A remedy would be to identify minimum requirements that must be satisfied by a measure of reliability, or more generally minimum requirements that must be satisfied by a measure of the condition (which may not be "reliability" in the intuitive sense) to be held equal for a fair assessment of the impact of coherence; and then show that strong Bayesian coherentism fails regardless of the measure of the condition we adopt. Until that is done, the game is not quite over yet for strong Bayesian coherentism.

As this example indicates, many remarkable results B&H obtain in this book are not necessarily final words on the subjects. The field is still young and the issues they address are worthy of further investigation. This small book contains powerful tools that probabilistic epistemologists can use in their own analyses of these subjects and beyond.[1]

Tomoji Shogenji
Philosophy Department
Rhode Island College
tshogenji@ric.edu