Bayesian Epistemology, by Luc Bovens and Stephan Hartmann, Oxford: Oxford University Press, 2003, Pp. ix + 159. H/b £37.50. P/b £16.99.

This short – but impressively rich – book breaks new ground in the application of advanced methods in Bayesian probabilistic modeling to various well known and important problems in contemporary epistemology and philosophy of science. This is, in many ways, a daring and impressive monograph which merits the careful scrutiny of analytic philosophers of all stripes. Owing to space limitations, I won't be able to delve into all of the fascinating and beautiful analyses contained in this book. Instead, my aim will be to give the reader a taste of the importance of the project, and of the impact that this nifty little book is already having on researchers in the field.

In the first two chapters, Bovens and Hartmann (B&H) engage in a very lucid and penetrating discussion of the notion of *coherence* as it appears in contemporary epistemology. Specifically, the views C.I. Lewis (as employed, *e.g.*, by Bonjour) are sympathetically articulated and precisely explicated using the machinery of modern Bayesianism. In Chapter one, the following two Lewisian principles concerning coherence and probability are analyzed in detail (throughout this review I will be paraphrasing so that I can stay within a reasonable word limit, and so as to best unify my own remarks – this will involve *inter alia* some changes of notation from B&H):

(L_1) Coherence is (*ceteris paribus*) *truth-conducive*. If one set of statements $\underline{S}$ is at least as coherent as another $\underline{S}'$ (*i.e.*, if $C(\underline{S}) \geq C(\underline{S}')$, where C is a measure of the *degree of coherence* of a set of statements) then $\underline{S}$ (*i.e.*, the conjunction of its elements) is at least as probable as $\underline{S}'$.

(L_2) Coherence orderings (*i.e.*, *complete* coherence orderings) over collections of sets of statements *supervene* on the *probabilistic* features of the sets involved. A (*complete*) coherence ordering over a collection of sets of statements **S** is fully determined by the probabilistic features of the sets of statements $\underline{S}$ contained in the collection **S**. [In other words, there exists a coherence measure C(•) that is defined *solely* in terms of *probabilistic* features of **S** such that, for all $\underline{S}$ and $\underline{S}'$ in **S**, $\underline{S}$ is at least as coherent as $\underline{S}'$ iff $C(\underline{S}) \geq C(\underline{S}')$, where this gives a *total* ordering over **S**.]

Chapter one contains a very important *impossibility theorem* (a la Arrow) to the effect that (L_1) and (L_2) *cannot* be jointly satisfied by *any* probabilistic measure $C(\underline{S})$ of the degree of coherence of a set of statements $\underline{S}$. This means that a Bayesian-Lewisian must give up either (L_1) or (L_2), or both. B&H's response to this crucial result of theirs is to give up (L_2), and to settle for a *quasi*-ordering (rather than a *complete* ordering) of coherence over collections of sets of statements. This seems like a very plausible strategy to me (*modulo* my worries below). But, I would also consider rejecting (L_1). After all, as B&H themselves point out, Lewis thought of coherence in terms of probabilistic *correlation*, and *not* (merely) in terms of *probability*. To wit:

(L_3) Coherence depends on *probabilistic correlation*. If one set of statement $\underline{S}$ is positively correlated, and another set $\underline{S}'$ is not positively correlated, then $\underline{S}$ is more coherent than $\underline{S}'$. More generally, for some suitable choice of measure of *degree of correlation* R(•), $C(\underline{S}) \geq C(\underline{S}')$ if and/or only if $R(\underline{S}) \geq R(\underline{S}')$. [Note: a set $\underline{S}$ is *positively correlated* *iff* each subset $\underline{s}$ of $\underline{S}$ is such that the probability of the conjunction of $\underline{s}$'s elements is greater than the product of the probabilities of its conjuncts. Measures of correlation are also called *confirmation* measures by contemporary Bayesians, who think of confirmation as degree of correlation.]

It seems to me (and I think B&H would agree) that (L_3) – at least its comparative clause – is something that Lewis would have accepted. And, *prima facie*, this principle is in tension with (L_1). As Carnap (in his 1962 book *Logical Foundations of Probability*, Chicago: University of Chicago Press) taught us, *truth*-conduciveness and *correlation*-conduciveness are not the same thing, and they sometimes come into conflict. Following Carnap (*op cit.*, new preface), we might think of (L_1) as *firmness*-conduciveness and (L_3) as *increase-in-firmness*-conduciveness. I prefer to follow modern Bayesian parlance, and so I will talk about (L_3) under the rubric "*confirmation*-conduciveness." While (L_1) and (L_2) cannot be simultaneously satisfied, perhaps (L_2) and (L_3) can. Indeed, work along just these confirmation-theoretic lines is now being done by Franz Dietrich and Luca Moretti ("On Coherent sets and the transmission of confirmation", unpublished manuscript, 2004, available through the *PhilSci Archive*), and the results are very promising. This is one of many examples illustrating the importance of the book, and how it is already shaping the way new research is being done in this area. The framework B&H develop in chapter one of this book is a model of how research in this area should be conducted.

In chapter two, B&H present their own probabilistic coherence measure, and they explain how it imposes (only) a *quasi*-coherence-ordering over collections of sets of statements. They also compare and contrast their measure with several other probabilistic coherence measures that have been proposed and defended in the recent literature (*e.g.*, those of Shogenji, Olsson, and Fitelson). This chapter is chock full of terrific examples, analyses, and intuition pumps. I won't be able to get into the details here. Instead, I want to return to (L_3) and the implicit role it plays in this chapter. The discussion of B&H suggests that all the measures they compare in chapter two satisfy (L_3). [Here, B&H seem to concur with Lewis's intuitions about (L_3), which explains their implicit assumption on this score.] Unfortunately, however, neither their coherence measure nor Olsson's satisfies (L_3). This has recently been shown by Wouter Meijs (in his "A Corrective to Bovens and Hartmann's Measure of Coherence", unpublished manuscript, 2004, available through the *PhilSci Archive*). I think it is a bit unusual to compare (L_3)-measures with non-(L_3)-measures in this way. Instead, I would say that measures like Olsson's and B&H's are *inadequate* (from a Lewisian point of view) insofar as they violate Lewis's (L_3). Happily, Meijs (*op. cit.*) shows how the B&H coherence measure can be repaired, so as to restore consistency with the Lewisian correlation principle (L_3). But, once the repair is made, some of the claims B&H seem to want to hold for their measure no longer do. Specifically, as Meijs (*op. cit.*) explains, when B&H are comparing their measure with Fitelson's in chapter two, they give a "counterexample" (pages 50–52) which is inconsistent with the conjunction of (L_3) and

(L_4) If all the statements in a set *S* are logically equivalent with each other, then the coherence of *S* is *maximal* (*i.e.*, no set can have a higher coherence than one whose statements are equivalent).

As far as I know, principle (L_4) is not one that was endorsed explicitly by Lewis. But, as Fitelson (in his 2003 paper "A Probabilistic Theory of Coherence", *Analysis* *63*: 194–199) explains, (L_4) is a plausible additional desideratum for any *confirmation*-based [*viz*, (L_3)-inspired] measure of coherence. Moreover, B&H *agree* with Fitelson that (L_4) should be adopted as a *desideratum* for measures of coherence. As far as I can tell, this all but neutralizes

B&H's criticism of Fitelson's measure of coherence, and it leaves them with a measure that is much closer to Fitelson's than one might have thought possible on one's first reading of B&H's chapter two. See Meijs (*op. cit.*) for further discussion of this important interplay between desiderata (L_3) and (L_4), and its consequences for constructing probabilistic measures of coherence. Once again, wherever one comes down on this particular dispute, one must gratefully acknowledge B&H's book for breaking new ground in our thinking about the relationship between coherence and probability. Before moving on from chapter two of B&H, I should mention another important line of thought that has recently spun off from the work of B&H on coherence. Consider the following principle:

(L_5) If two *sets* of statements are logically equivalent (*i.e.*, if they have precisely the same logical consequences), then either they are both coherent or they are both incoherent.

Interestingly, (L_5) and (L_4) cannot be jointly satisfied by *any* non-trivial coherence measure, since (L_5) and (L_4) together entail that either *all* sets are coherent or *all* sets are incoherent! This, and many other fascinating (if not counterintuitive) recent results have been proved and explained by Luca Moretti and Ken Akiba (in their "Probabilistic Measures of Coherence and the Problem of Belief Individuation", unpublished manuscript, 2004, available from the *PhilSci Archive*). It is probably not unfair to say that none of this recent work on probabilistic (Lewisian) approaches to coherence would have appeared if B&H hadn't embarked on the pioneering project which is articulated in the early chapters of this book.

In chapter three, B&H show how one can apply the powerful machinery of the theory of Bayesian Networks to provide formal explications the epistemologically central notion of *reliability* (of, say, witnesses, or measuring instruments, or belief generating processes *etc*.). Bayesian Networks are, essentially, efficient (and, if presented properly, visually illuminating) ways of encoding probability distributions exhibiting certain kinds of conditional independence relations. These models began to appear *en masse* in various technical monographs in the 1980's (most notably, in Judea Pearl's (1988) masterpiece *Probabilistic Reasoning in Intelligent Systems*, San Francisco: Morgan Kaufmann). Subsequent applications of Bayesian Network methodology in the *statistical* sciences (especially, in the areas of knowledge discovery and data mining) have become ubiquitous. Moreover, applications of Bayes Nets in cognitive science (especially, in the area of *causal learning*) have also taken off [see Glymour's (2001), *The Mind's Arrows*, Cambridge, Mass: MIT Press]. Various *philosophical* applications of Bayesian Networks have appeared in the literature since then as well. But, these have been mainly restricted to the context of probabilistically informed accounts of *causation* or *causal inference* [*e.g.*, as in the work of Glymour *et al.* on the TETRAD project]. In chapter three, B&H begin to show how fruitful applications of Bayesian Networks can be in *scientific inference* and *epistemology* more generally. This chapter contains one of the most accessible, general introductions to Bayesian Networks I have seen (one well suited to philosophy graduate students and researchers). And, at the same time, it explains how Bayesian Network *models* can be used to aid in the understanding of various kinds of inferences involving multiple sources of information with varying degrees of reliability. Novel and insightful applications to *jury voting* and the *conjunction fallacy* (a la Tversky and Kahneman) are presented as nice, self-contained illustrations of the method. The remainder of the book elaborates on other, more sophisticated and powerful applications of Bayesian Network methodology in epistemology and the philosophy of science.

In chapter four, B&H show how the simple ideas from chapter three about modeling multiple sources of information with varying degrees of reliability can be applied to several key problems in confirmation theory. Specifically, B&H construct simple and intuitive Bayesian network models of various kinds of experimental set-ups. While these models are idealizations, to be sure, they are quite powerful and useful when applied to traditional problems in confirmation theory. Again, space limitations prevent me from getting into any of the details of their models. But, they provide detailed and rigorous new analyses of both the *variety of evidence* phenomenon and the *Duhem-Quine problem*. It is part of folk philosophy of science (dating back at least to Nagel) that more "various" or "diverse" evidence should (*ceteris paribus*) confirm more strongly than more "narrow" (or less "diverse") evidence does. Several Bayesian authors (including Horwich and others) have tried to provide probabilistic explications of this phenomenon, as we think it occurs in actual hypothesis testing in science. The approach of B&H is completely novel and quite ingenious. Their Bayesian Network methodology provides many new confirmation-theoretic insights here that I think would have been otherwise inaccessible to traditional Bayesian confirmation theoretic analyses. Indeed, some of the things Bayesian confirmation theorists have taken for granted in this context can now be (at least to some extent) *explained* using the methods B&H introduce in this chapter. In this sense, I think their approach complements the traditional Bayesian approaches nicely (rather than being in competition with them, as some have argued). The Duhem-Quine problem is another oldie-but-goodie that B&H take on in this chapter of the book. In the 1910's, Duhem noted that no theory *in isolation* entails any observational consequences. Rather, only in conjunction with *auxiliary hypotheses* do theories entail predictions. Later, Quine picked-up on this Duhemian theme (and in classic, *holistic* Quinean fashion) reasoned that Duhem's remark implies that we can never be sure which *part* of our theory + auxiliaries is to blame for a false prediction. That is, if *T* & *A* entails not-*E*, and *E* is observed, we cannot tell which of *T* or *A* (or both) is to blame for the incorrect prediction (we don't know where to "aim the *modus tollens*''). Bayesians have tried various approaches (with varying degrees of sophistication and success) to try to explain when *T* (or *A*) is *differentially disconfirmed* by *E* in such cases (*e.g.*, when *T* is disconfirmed *less strongly* by *E* than *A* is). But, these approaches tend to make rather strong (and often historically implausible) assumptions about the probabilistic relationships between *T*, *A*, and *E*. Specifically, traditional Bayesian approaches typically assume that *T* and *A* are *probabilistically independent*. While this (strong) assumption allows Bayesians to prove some salient theorems about Duhem-Quine scenarios, it would be nice to see an approach that operated with a weaker and more plausible set of assumptions. In this chapter, B&H show how Bayesian Network models of Duhem-Quine situations can do just that. Their models seem (to me) to be much more scientifically (and historically) plausible, and their conclusions are quite interesting and refreshing. The models of B&H allow us to weaken the traditional independence assumptions without sacrificing the probative value of the resulting models. And, in B&H's account, the reliability of the instruments (experimental apparatus) used in the test of the theory + auxiliaries plays a crucial role in their response to Duhem and Quine. While the devil is in the details (and space is too short to delve into them here), one thing is clear: the Bayesian Network models introduced in this chapter by B&H have a lot to offer contemporary practitioners of Bayesian confirmation theory.

Chapter five is the final chapter of the book, and, in many ways, it is bound to be the most interesting to many readers. Chapter five is all about *testimony*. The central question of this chapter is: how should we model testimonial inputs from multiple witnesses with varying degrees of independence and varying degrees of reliability? This is one of the oldest set of questions for probabilistic approaches to epistemology. As Earman [in his (2000) book *Hume's Abject Failure*, Oxford: Oxford University Press] skillfully explains, questions about the probabilistic aggregation of testimonial evidence have been around in philosophy for hundreds of years. Hume and Price (among others) made extensive use of arguments of this kind, especially, in the context of arguments concerning "miracles". In chapter five, B&H aren't so interested in miracle arguments *in particular* (although, their models are bound to work quite well for such purposes!). Rather, as is the theme throughout the book, they are after a general *framework* that utilizes Bayesian Network modeling techniques to provide new theoretical insights into the general class of problems of which Hume's and Price's (and Earman's) are special cases. Here, they show their usual theoretical aplomb, and they develop several new, illuminating approaches to this important class of epistemological problems. Amusingly, their ultimate application of these ideas is to the problem of "shopping for consumer products". Here, as elsewhere in the book, I learned several deep lessons about probability and evidence. I am sure that I will use some of the methods developed in this chapter (and the rest of the book) in my own research. The possibilities for developing all sorts of models along these lines are seemingly limitless.

Bovens and Hartmann have much to be proud of here. This book is just the tip of what promises to be a rather large iceberg in the growing field of "probabilistic epistemology". And, B&H are two of its most energetic and productive pioneers. The theoretical and philosophical insights contained in this little book are belied by its brevity. While the book is often dense with subtle technical and philosophical arguments, it more than repays a minute reading. I can't think of another book this short that taught me this much. Anyone who is interested in the relationship between probability, evidence, and knowledge, and/or anyone who wants to see the state of the art in the application of Bayesian modeling techniques to the philosophy of science should read this book straightaway.

Department of Philosophy
University of California
Berkeley, CA 94720–2390
USA

BRANDEN FITELSON