an apple before me' at *t whether or not S believes P at t*. Since P is non-inferentially, propositionally justified for S at *t*, P turns out to be an item of S's evidence at *t* on E = NPJ. Goldman argues through examples that E = NPJ is capable of explaining all that E = K explains and additional things. He concludes that E = NPJ is rationally preferable to E = K. Williamson's response (pp. 308–12) systematically questions Goldman's examples and arguments. As always Williamson's observations are subtle and interesting, but I doubt that many readers will be persuaded. There is an elementary reason why E = NPJ looks prima facie more plausible than E = K. E = NPJ can straightforwardly account for very basic platitudes such as the commonplace view that S can have *misleading* evidence (P can be evidence for S on E = NPJ even if P is false) or *overlooked* evidence (P can be evidence for S on E = NPJ even if S does not believe P). In contrast E = K cannot easily account for these basic platitudes. For since knowledge entails truth and belief, any piece of evidence that S possesses, on E = K, must be true and believed by S.

I have been able to survey only a very few of the numerous and stimulating criticisms of knowledge-first epistemology that the reader can find in *Williamson on Knowledge*. Williamson's ground-breaking epistemological position might appear prima facie plausible because of the discouraging list of failed attempts to provide a reductive analysis of knowledge. Yet its real strength depends on the number of coherent and mutually supporting theses that follow from Williamson's central assumption that knowledge is basic and unanalysable. The appropriate method of appraisal of knowledge-first epistemology can only consist in the assessment of many or most of these consequences at once. The essays in *Williamson on Knowledge* jointly attempt at such an overall evaluation. This is why this volume turns out to be so intriguing and valuable.

Department of Philosophy LUCA MORETTI
University of Aberdeen
Old Brewery
High St
Aberdeen AB24 3UB
Scotland
UK
doi:10.1093/mind/fzt005 Advance Access publication 21 March 2013

Reliable Reasoning, by Gilbert Harman and Sanjeev Kulkarni. Cambridge MA: MIT Press, 2007. Pp. 108.

The aim of the book is to give a non-technical introduction to statistical learning theory at undergraduate level. Statistical learning theory is concerned with the reliability of rules for classifying a new case — e.g. diagnosing a disease in a new patient — on the basis of other features of the case and a large stock of past cases and their features and classifications. The book is based on a course on learning theory and epistemology given to undergraduate students in electrical engineering and in philosophy at Princeton.

It is a short book, with four chapters. The first chapter is on 'the problem of induction'. Note though, that the book is not concerned with what philosophers normally take to be '*the* problem of induction' — the problem of justifying induction — but rather with the problem of assessing the reliability of inductive rules. The chapter then sets out to debunk the commonly held view that there are two sorts of reasoning, deductive and inductive, and two sorts of arguments, deductive and inductive, by means of arguments to which I will return below. Chapter two introduces the idea of using enumerative induction to learn classification rules, and for estimating the values of continuous variables. It introduces the *VC-dimension* of a set of classification rules, which is perhaps the most important concept in statistical learning theory. Chapter three discusses induction rules which work by trying to rank hypotheses by their simplicity. Chapter four discusses applications of statistical learning theory to neural networks and support vector machines in machine learning.

The book does a good job at presenting the main ideas underlying statistical learning theory. Members of our interdisciplinary Centre for Reasoning at the University of Kent read the book in a reading group and found that those who were unfamiliar with statistical learning theory were able to grasp the central intuitions, while those who were familiar with statistical learning theory had some controversial philosophical claims to get their teeth into. Philosophical topics include the nature of induction, Goodman's new problem of induction, simplicity, mental processes for reasoning, and moral particularism. Note that the book is intentionally light on detail, so it may help to have someone familiar with both the statistical and the philosophical details on hand to fill in gaps and answer students' questions. This restricts the applicability of the book somewhat.

While the book largely succeeds in its aim of presenting the rudiments of statistical learning theory in an accessible way, it does so at the expense of presenting a questionable view of the key concept under consideration, namely that of induction. As mentioned above, the book sets up the problem of the reliability of induction by trying to debunk the usual view that there is a distinction between deductive and inductive reasoning and a distinction between deductive and inductive arguments. The usual view has it that in a valid deductive argument the truth of the premisses $\varphi_1, \ldots, \varphi_n$ forces the truth of the conclusion ψ (we write $\varphi_1, \ldots, \varphi_n \models \psi$), but that in an inductive argument the premisses merely make the conclusion more or less plausible; we can write $\varphi_1^{X_1}, \ldots, \varphi_1^{X_n} \mathrel{|}\!\approx \psi^Y$ to say that the premiss propositions $\varphi_1, \ldots, \varphi_n$ together with their respective levels of plausibility $X_1, \ldots, X_n$ confer some level Y of plausibility on the conclusion proposition ψ (See section 1.1 of R.

Haenni, J.-W. Romeijn, G. Wheeler, and J. Williamson, *Probabilistic Logics and Probabilistic Networks*, Synthese Library, Dordrecht: Springer, 2011). Various deductive and inductive logics have been put forward to formalize such arguments. The usual view has it that while these logics may have applications that do not directly concern reasoning (e.g. to circuit design in computer science), reasoning provides a key application (and often *the* leading application) in that many logics are developed with the application to reasoning in mind. An instance of reasoning that corresponds closely to a deductive argument is an instance of deductive reasoning, while an instance that corresponds to an inductive argument is an instance of inductive reasoning.

Harman and Kulkarni take issue with this standard view, offering two main objections. The first objection hinges on the relationship between arguments and reasoning. The authors quite rightly note that reasoning can lead to the abandoning of something that one started out believing, but they claim that 'an obvious difficulty with the traditional picture is its implication that reasoning is always a matter of inferring new things from what one starts out believing' (p. 5). Of course this would be a problem were it an implication of the traditional picture, but it is not. A deductive entailment relationship $\varphi_1, \ldots, \varphi_n \vDash \psi$ provides good reason not to grant all of $\varphi_1, \ldots, \varphi_n, \neg\psi$, but it does not imply that one should continue to grant $\varphi_1, \ldots, \varphi_n$ if one started out granting those propositions. The fault here seems to lie not with the traditional picture but rather with the spurious implication drawn by the authors. A more sensible claim might take something like the following form: if (i) $\varphi_1, \ldots, \varphi_n \vDash \psi$, (ii) $\varphi_1, \ldots, \varphi_n$ are consistent, (iii) ψ is relevant to your current context of inquiry, and (iv) you grant *just* that $\varphi_1, \ldots, \varphi_n$ and that ψ follows deductively from $\varphi_1, \ldots, \varphi_n$, then you should infer ψ.

Their second objection to the standard picture concerns the relationship between deduction and induction. Harman and Kulkarni claim:

> It is a category mistake to treat deduction and induction as belonging to the same category. Deductive arguments are abstract structures of propositions, whereas inductive reasoning is a process of change in view. … There is deductive logic, but it is a category mistake to speak of inductive logic. (pp. 7–8)

The motivation behind such a stance is not entirely clear from the text, but it seems to rest again on the connection between arguments and reasoning:

> One sort of theory of probability is an abstract mathematical subject. How it is to be applied to reasoning is not part of the mathematics. The same point holds for decision theories that appeal to utilities as well as probabilities. These theories offer extended accounts of consistency or 'coherence' of belief but leave open in what way such consistency or coherence is relevant to reasoning. (p. 9)

It is certainly true that rules for applying a formalism are seldom a part of the formalism itself. But this platitude does not provide grounds for questioning the standard view of the relationship between induction and deduction.

Indeed, something that inductive and deductive logics have in common is that they are often pitched at a formal level, with questions associated with their application to reasoning left to the 'knowledge engineer' rather than the logician to sort out, as the contextual details of the particular application may have a crucial bearing on the way in which the logic is applied. Nevertheless, the application to reasoning often plays a leading role in the development of the logic, and that is true as much for Bayesian inductive logics (to which Harman and Kulkarni appear to be alluding in the above quote) as for deductive logics and the plethora of non-monotonic logics. Indeed, Bayesian inductive logics are concerned primarily with degrees of *belief*, with an inductive entailment relationship of the form $\varphi_1^{X_1}, \ldots, \varphi_1^{X_n} \approx \psi^Y$ being read as: if $P(\varphi_1) \in X_1, \ldots, P(\varphi_n) \in X_n$ capture all the given constraints on rational degree of belief, then $P(\psi) \in Y$ is a derived constraint on rational degree of belief. If anything, this sort of inductive logic wears its leading application on its sleeve — more so than does deductive logic.

In short, neither objection to the standard view of deduction and induction bears much scrutiny and the book fails to provide convincing grounds for rejecting this view.

While the preliminary discussion of induction will, I think, rather confuse and mislead students, subsequent chapters of the book present the core ideas of statistical learning theory in an engaging way, and chapters two to four would provide valuable background reading for part of a course on induction. There is a lot of exciting formal work done on machine learning that is relevant to epistemology and normative reasoning (see, for example, D. Corfield, 'Varieties of Justification in Machine Learning', *Minds and Machines,* 20 (2010), pp. 291–301); the book is to be commended for making some of this work more accessible to philosophy students.

(Work on this review was supported by a grant from the UK Arts and Humanities Research Council.)

Department of Philosophy JON WILLIAMSON
SECL
University of Kent
Canterbury CT2 7NF
UK
j.williamson@kent.ac.uk
doi:10.1093/mind/fzt006 Advance Access publication 4 April 2013

