Journal of Applied Logic 11 (2013) 468–486

Contents lists available at SciVerse ScienceDirect

Journal of Applied Logic

www.elsevier.com/locate/jal

From Bayesian epistemology to inductive logic

Jon Williamson

University of Kent, Canterbury, United Kingdom

ARTICLE INFO

Article history:
Available online 18 March 2013

Keywords:
Bayesianism
Objective Bayesianism
Bayesian epistemology
Maximum entropy
Maxent
Inductive logic
Probabilistic logic
Probability logic

ABSTRACT

Inductive logic admits a variety of semantics (Haenni et al. (2011) [7, Part 1]). This paper develops semantics based on the norms of Bayesian epistemology (Williamson, 2010 [16, Chapter 7]). Section 1 introduces the semantics and then, in Section 2, the paper explores methods for drawing inferences in the resulting logic and compares the methods of this paper with the methods of Barnett and Paris (2008) [2]. Section 3 then evaluates this Bayesian inductive logic in the light of four traditional critiques of inductive logic, arguing (i) that it is language independent in a key sense, (ii) that it admits connections with the Principle of Indifference but these connections do not lead to paradox, (iii) that it can capture the phenomenon of learning from experience, and (iv) that while the logic advocates scepticism with regard to some universal hypotheses, such scepticism is not problematic from the point of view of scientific theorising.

1. Bayesian epistemology as semantics for inductive logic

This section introduces the use of Bayesian epistemology as semantics for inductive logic. The material presented here is based on Williamson [16], to which the reader is referred for more details.

1.1. Bayesian epistemology: A primer

At root, Bayesian epistemology concerns the question of how strongly one should believe the various propositions that one can express. The Bayesian theory that answers this question can be developed in a number of ways, but it is usual to base the theory on the betting interpretation of degrees of belief. According to the betting interpretation, one believes proposition θ to degree x iff, were one to offer a *betting quotient* for θ—a number q such that one would pay qS to receive S in return should θ turn out to be true, where unknown stake $S \in \mathbb{R}$ may depend on q—then $q = x$.

This interpretation of degrees of belief naturally goes hand in hand with the claim that, were one to bet according to one's degrees of belief via the above betting set-up, then one shouldn't expose oneself to avoidable losses. In particular, arguably one's degrees of belief should minimise worst-case expected loss.

This starting point—the betting interpretation together with the loss-avoidance claim—can then be used to motivate various rational norms that answer the main question facing Bayesian epistemology, i.e., that specify how strongly one should believe the various propositions that one can express. Since the context of this paper is inductive logic, we will be particularly concerned with propositions expressed in a logical language. Suppose then that $\mathcal{L}_n$ is a propositional language on elementary propositions $A_1, \ldots, A_n$, with $S\mathcal{L}_n$ the set of propositions formed by recursively applying the usual logical connectives to the elementary propositions. Let Ω_n be the set of *atomic states* of $\mathcal{L}_n$, i.e., propositions ω_n of the form $\pm A_1 \wedge \cdots \wedge \pm A_n$, where $\pm A_i$ is either A_i or its negation.

E-mail address: j.williamson@kent.ac.uk.
URL: http://www.kent.ac.uk/secl/philosophy/jw/.

1570-8683/$ – see front matter
http://dx.doi.org/10.1016/j.jal.2013.03.006

One norm of rational belief, which we shall call the *Probability Norm*, says that one's degrees of belief should satisfy the axioms of probability, for otherwise, in the worst case, stakes are chosen that ensure positive expected loss—equivalently, stakes are chosen that ensure positive loss whichever atomic state turns out to be true (a so-called *Dutch book*). Thus degrees of belief should be probabilities in order to minimise worst-case expected loss:

Theorem 1. *Define function $P : S\mathcal{L}_n \to \mathbb{R}$ by $P(\theta) =$ a given agent's betting quotient for θ. The agent's bets on propositions expressible in $\mathcal{L}_n$ avoid the possibility of a Dutch book if and only if they satisfy the axioms of probability*:

P1. $P(\omega_n) \geqslant 0$ *for each* $\omega_n \in \Omega_n$,
P2. $P(\tau) = 1$ *for some tautology* $\tau \in S\mathcal{L}$,
P3. $P(\theta) = \sum_{\omega_n \models \theta} P(\omega_n)$ *for each* $\theta \in S\mathcal{L}$.

See Williamson [16, Theorem 3.2] for a proof. This is known as the Dutch Book Theorem, or the Ramsey–de-Finetti Theorem. P1–3 offer one way of expressing the axioms of probability over the propositional language $\mathcal{L}_n$: this axiomatisation makes it clear that a probability function is determined by its values on the atomic states of $\mathcal{L}_n$.

A second norm, the *Calibration Norm* or *Principal Principle*, says that one's degrees of belief should be calibrated to known physical probabilities. In particular, if one knows just the physical probability $P^*(\theta)$ of θ then one should set one's betting quotient for θ to this physical probability, $P(\theta) = P^*(\theta)$, for otherwise in the worst case stakes will be chosen that render the expected loss positive.[1] While in the case of the Probability Norm, minimising worst-case expected loss is equivalent to avoiding sure loss (a Dutch book), in this case the link appears in a long run of bets rather than in the single bet on θ itself: if one were repeatedly to bet on θ-like events with the same betting quotient for each bet, then one would be susceptible to sure loss in the long run [16, pp. 39–42].

Of course in general one might know only of certain constraints on physical probability, rather than individual physical probabilities such as $P^*(\theta)$. In such a case the Calibration Norm would say that if one knows that the physical probability function P^* lies within some set $\mathbb{P}^*$ of probability functions, then one's belief function P should lie within the convex hull $\langle\mathbb{P}^*\rangle$ of this set of probability functions. The reason being that if one remains in the convex hull one's bets do not have demonstrably greater than the minimum worst-case expected loss (equivalently, one can't be forced to lose money in the long run), but outside the convex hull one can be sure of sub-optimal expected loss (equivalently, one can be forced to lose money in the long run).

A third norm, the *Equivocation Norm*, says that one should not adopt extreme degrees of belief unless forced to by the Probability Norm or the Calibration Norm: one's degrees of belief should be equivocate sufficiently between the basic possibilities that one can express (i.e., between the atomic states of $\mathcal{L}_n$). Equivalently, one's belief function P should satisfy the constraints imposed by the other norms and otherwise should be sufficiently close to the *equivocator* function that gives the same probability to each of the 2^n atomic states, $P_=(\omega_n) = 1/2^n$ for each $\omega_n \in \Omega_n$. (Distance between probability functions is measured by Kullback–Leibler divergence, $d(P, Q) = \sum_{\omega_n \in \Omega_n} P(\omega_n) \log P(\omega_n)/Q(\omega_n)$.) Again this norm can be justified by minimising worst-case expected loss; the argument goes as follows [16, pp. 63–65]. In the absence of knowledge about one's losses, one should take the loss function to be logarithmic by default, i.e., one should assume that one will lose $-\log P(\omega_n)$ where ω_n is the atomic state that turns out true, for such a loss function is the only one that satisfies various desiderata that are natural to impose on a default loss function. But then, under some rather general conditions, the P that satisfies constraints imposed by other norms but minimises worst-case expected loss (the *robust Bayes* choice of P) is the P that is closest to the equivocator.

Note that we need the qualification that degrees of belief should be 'sufficiently' close to the equivocator in order to handle the case where there is no function closest to the equivocator. For example, one might know that a coin is biased in favour of heads, so that $P^*(H) > 1/2$ where H signifies heads at the next toss. Arguably then by the Calibration Norm one ought to believe H to degree greater than 1/2, $P(H) > 1/2$. But there is no degree of belief greater than 1/2 that is closest to $P_=(H) = 1/2$. Therefore the most one can expect of an agent is that $P(H)$ is sufficiently close to 1/2, where what counts as sufficiently close depends on pragmatic considerations such as the required accuracy of calculations.

In sum, the betting interpretation together with the loss-avoidance claim lead to three norms: Probability, Calibration and Equivocation. Bayesian epistemologists disagree as to whether to endorse all these norms. All accept the Probability Norm, most accept some version of Calibration Norm, but few accept the Equivocation Norm. The Probability Norm on its own leads to what is sometimes called *strict subjectivism*, Probability together with Calibration leads to *empirically-based subjectivism* and all three norms taken together lead to *objectivism*. Note that strict subjectivism doesn't imply that determining appropriate degrees of belief is entirely a question of subjective choice, inasmuch as the Probability Norm

[1] For simplicity of exposition we take physical probability P^* to be single-case here, defined over the propositions of the agent's language $\mathcal{L}$. But we could instead take physical probability to be generic, defined over repeatedly instantiatable outcomes, and then consider in place of P^* the single-case consequences of physical probability, i.e., the constraints that known generic physical probabilities impose on the agent's degree of belief. See Williamson [19] on this point.