Brit. J. Phil. Sci. **59** (2008), 201–211

Bayesian Confirmation by Uncertain Evidence: A Reply to Huber [2005]

Vincenzo Crupi, Roberto Festa and Tommaso Mastropasqua

ABSTRACT

Bayesian epistemology postulates a probabilistic analysis of many sorts of ordinary and scientific reasoning. Huber ([2005]) has provided a novel criticism of Bayesianism, whose core argument involves a challenging issue: confirmation by *uncertain* evidence. In this paper, we argue that under a properly defined Bayesian account of confirmation by uncertain evidence, Huber's criticism fails. By contrast, our discussion will highlight what we take as some new and appealing features of Bayesian confirmation theory.

1 Introduction

Bayesian epistemology postulates a probabilistic analysis of many sorts of ordinary and scientific reasoning. Also, contemporary Bayesians typically endorse a subjective reading of probability, i.e. interpret probabilities as degrees of subjective belief. Huber ([2005]) has provided a novel criticism of Bayesianism, whose core argument involves a challenging issue: confirmation by *uncertain* evidence, i.e. evidence which has *not* been ascertained. In order to assess Huber's argument, it is crucial to combine Bayesian confirmation theory with Jeffrey conditionalisation. In this paper, we will argue that, when properly merged with Jeffrey conditionalisation, Bayesian confirmation theory escapes Huber's criticism and yields some new and appealing results.

Our discussion will proceed as follows. First, we will outline a generalised version of Bayesian confirmation theory which can be readily applied under

doi:10.1093/bjps/axn008

Jeffrey conditionalisation. Then, we will review a crucial requirement at the core of Huber's argument and show that it is equivocal. We will argue that on one reading it amounts to a compelling principle, whereas on an alternative reading it turns out to be highly implausible. Finally, we will show that our account of Bayesian confirmation by uncertain evidence appropriately captures the former version of the requirement and violates the latter.

2 Uncertain Evidence and Bayesian Confirmation

For our purposes, we will consider a (non-empty) set of statements Γ closed under truth-functional operators, such as negation, conjunction and disjunction. Bayesians commonly assume that, at a given time x, the belief state of an agent A concerning the statements in Γ is represented by a probability function P_x defined over that set.

It may occur that, from time x to y, A experiences a change in opinion concerning a particular $e \in \Gamma$ (provided that $P_x(e)$ is not extreme, i.e. $0 < P_x(e) < 1$). One important question is, then: how should A's beliefs in other statements belonging to Γ change as a consequence?

Up to the mid-sixties, Bayesians had a ready answer only for the special case in which, at time y, A has come to believe that e is *certainly true*, so that $P_y(e) = 1$ (and, correspondingly, $P_y(\neg e) = 0$). 'Classical' Bayesian updating or conditionalisation (BC) postulates that

$$\text{(BC) If } P_y(e) = 1, \text{ than for any } h \in \Gamma,\ \ P_y(h) = P_x(h \mid e).$$

However, it may surely also occur that A's degree of belief in e changes from time x to y *without* reaching certainty. What will be the value of $P_y(h)$ *then*? Richard Jeffrey has suggested a natural and elegant way to generalise classical Bayesian conditionalisation ([1965], Chapter 11; also see Jeffrey [2004], pp. 53–5). In Jeffrey conditionalisation (JC), it is assumed that

$$\text{(JC) For any } h \in \Gamma,\ \ P_y(h) = P_x(h \mid e)P_y(e) + P_x(h \mid \neg e)P_y(\neg e).$$

Thus in (JC), $P_y(h)$ is computed as an average of the 'old' conditional probabilities of h on e versus $\neg e$, weighted by the current probabilities of e and $\neg e$, respectively. It is easy to see that Jeffrey conditionalisation is a proper generalisation of classical Bayesian updating in the sense that (JC) implies (BC) (not the converse). Under Jeffrey conditionalisation, however, a change in belief about e prompts the updating of the prior probability $P_x(h)$ to a new value $P_y(h)$ which is generally *not* identical to either the conditional $P_x(h \mid e)$ (except when e does become certainly true) or the conditional $P_x(h \mid \neg e)$ (except when e becomes certainly false), but rather lies between those two values.

Now consider the Bayesian notion of *confirmation*. Bayesian confirmation theory has been commonly elaborated and applied on the background of classical Bayesian updating. The issue has been to formalise the impact on a hypothesis h of the (often implicit, and quite restrictive) assumption of evidence e having been *ascertained*, i.e. precisely in case $P_x(e) \neq P_y(e) = 1$. Then h is said to be confirmed if and only if $P_y(h) = P_x(h \mid e) > P_x(h)$ and to be disconfirmed if and only if $P_y(h) = P_x(h \mid e) < P_x(h)$. (If $P_y(h) = P_x(h \mid e) = P_x(h)$, it is said that coming to know that e is *neutral* for h.)

Can Bayesian confirmation theory be extended to cases such that from time x to y the probability of e changes, but the assumption of e having been ascertained at y is relaxed? In other terms, is there any natural way to parallel Jeffrey's generalisation of classical Bayesian updating in the framework of confirmation, and provide a plausible probabilistic account of confirmation by *uncertain* evidence? In what follows we'll claim that the answer is in the positive. (In essence, we will be following a proposal already made in Festa [1999], pp. 56–9.)

It is well known that various alternative measures of confirmation have been proposed and defended by Bayesian theorists (see Festa [1999] and Fitelson [1999]). For our present purposes, it will be convenient to focus on a core set of such confirmation measures which share the following interesting property: they can be defined by means of a function f depending *only* on $P(h \mid e)$ and $P(h)$, f being a strictly increasing function of the former value and a non-increasing function of the latter. We will call such confirmation measures 'classically P-incremental'. Classically P-incremental measures include

- the 'difference' measure, first defined by Carnap ([1950/1962], p. 361) as:

$$D(h, e) = P(h \mid e) - P(h);$$

- the 'ratio' measure, first defined by Keynes ([1921], pp. 150–5) as:

$$R(h, e) = \frac{P(h \mid e)}{P(h)};$$

- the 'odds ratio' measure, first conceived by Alan Turing (as reported by Good [1950], pp. 62–3) as:

$$OR(h, e) = \frac{P(h \mid e)/P(\neg h \mid e)}{P(h)/P(\neg h)};$$

- and the following measure, recently discussed by Crupi et al. [2007][1]

$$Z(h,e) = \begin{cases} \frac{P(h|e)-P(h)}{1-P(h)} & \text{if } P(h \mid e) \geq P(h) \\ \frac{P(h|e)-P(h)}{P(h)} & \text{otherwise.} \end{cases}$$

Notice that, in the notation adopted here, $P(h) = P_x(h)$ whereas, under classical Bayesian conditionalisation, $P(h \mid e) = P_y(h)$. Thus, when classical Bayesian conditionalisation applies, the above definitions can immediately be converted into

$$\begin{aligned} D_{x,y}(h) &= P_y(h) - P_x(h), \\ R_{x,y}(h) &= \frac{P_y(h)}{P_x(h)}, \\ OR_{x,y}(h) &= \frac{P_y(h)/P_y(\neg h)}{P_x(h)/P_x(\neg h)}, \end{aligned}$$

$$Z_{x,y}(h) = \begin{cases} \frac{P_y(h)-P_x(h)}{1-P_x(h)} & \text{if } P_y(h) \geq P_x(h) \\ \frac{P_y(h)-P_x(h)}{P_x(h)} & \text{otherwise.} \end{cases}$$

Here, the double subscript 'x,y' highlights the fact that confirmation is *relative* in an important sense: it is crucial for confirmation (disconfirmation) of a hypothesis h by a change in opinion about e to occur in the *shift* from one probability distribution, P_x, to another, P_y, such that $P_x(e) \neq P_y(e)$.

But now our claim is that these latter formulas *already represent* straightforward ways to generalise the corresponding confirmation measures as usually defined in the literature. This is because $D_{x,y}(h)$, $R_{x,y}(h)$, $OR_{x,y}(h)$ and $Z_{x,y}(h)$ all measure (although in different ways) the departure from the initial probability of h—$P_x(h)$—of an appropriately updated probability $P_y(h)$. Under Jeffrey conditionalisation, generalised confirmation will amount to the departure from prior probability *not* of the conditional $P_x(h \mid e)$ (which, again, is not attained except in the special case of e having in fact being ascertained), but rather of the updated probability $P_y(h)$ to which a change in belief about the uncertainty of e will lead. Clearly, for *any* classically P-incremental Bayesian confirmation measure, a generalised version can be devised along these lines. Importantly, by such a move, any classically P-incremental measure will also

[1] Advocates of measure D include Eells ([1982]) and Earman ([1992]). Advocates of measure R include Horwich ([1982]) and Milne ([1996]). Advocates of measure OR include Good himself ([1950], [1983]) as well as Fitelson ([2001]). To the best of our knowledge, the earliest appearance of Z is in a rather fleeting discussion in Rescher ([1958], p. 87), where a different (not P-incremental) confirmation measure is ultimately endorsed. The positive branch of Z is identical to Rips's ([2001], p. 129) quantitative measure of 'inductive strength' and ordinally equivalent to a confirmation measure proposed by Gaifman ([1979], p. 120). Further occurrences are in the literature on expert systems and expert judgment (see Cooke [1991], p. 57).

satisfy a *generalised* condition of P-incrementality, i.e. it will be expressible by means of a function f depending only on $P_y(h)$ and $P_x(h)$, f being a strictly increasing function of the former value and a non-increasing function of the latter.[2]

As a final remark, notice that, even beyond Jeffrey conditionalisation, generalised P-incremental measures are suitable of application under any kind of updating rule considered in probability kinematics. This is because they only require defined values of $P_x(h)$ and $P_y(h)$ themselves, however related.[3]

3 Bayesian Confirmation by Uncertain Evidence: Test Cases and Basic Principles

Huber ([2005]) has provided a nice hypothetical test case for Bayesian confirmation by uncertain evidence. Suppose

$h =$ all Scots wear kilts,

$e =$ the Scottish guy Stephen wears a kilt.

Notice that $h \models e$ (not the converse), so that the probability of the latter given the former must equal 1.[4] Also, a Bayesian account would provide an agent A with initial probabilities $P_x(e)$ and $P_x(h)$ such that $P_x(e) > P_x(h)$, again because of the logical relationship between the two statements. It is then assumed that A is initially uncertain about both e and h, so that both $P_x(e)$ and $P_x(h)$ are not extreme. It follows that coming to believe with certainty that e would confirm h, i.e. $P_x(h\,|\,e) > P_x(h)$.

Suppose that A, who is not wearing her glasses, looks at Stephen and comes to subjectively believe that e with a moderate level of confidence, assumed to be represented by

$$P_y(e) = 0.6$$

Importantly, Huber's ([2005]) discussion of the example clearly suggests that $P_y(e) > P_x(e)$, i.e. that A's observation has increased her confidence in e.

[2] Such a generalised P-incrementality condition will play an important role in what follows. For this reason, we are leaving aside here various confirmation measures proposed by Bayesian theorists which are demonstrably not P-incremental: see (Carnap [1950/1962], p. 360; Nozick [1981], p. 252; Mortimer [1988], Section 11.1; Christensen [1999], p. 449; Joyce [1999], Chapter 6).

[3] Over the years, Bayesian theorists dealing with probability kinematics have considered various forms of updating, as prompted by different kinds of information (see, for instance, van Fraassen [1980]; Jeffrey [1992], Chapters 6–7).

[4] Strictly speaking, in order to have $h \models e$, we should include 'Stephen is Scottish' as a separate background knowledge statement within Γ and modify notation accordingly. We embedded it in e simply for ease of exposition. This, however, has no effect on the issue discussed in this paper.

Now consider A looking at Stephen *with her glasses on* and coming to subjectively believe that e with a *high* level of confidence, such that

$$P_z(e) = 0.9.$$

Commenting on his example, Huber remarks that 'if some e speaks in favour of some h—say, because it is a logical consequence of the latter—then [...] getting to know that e is probably true should provide confirmation for h—and *the more probable it is that* e *is true, the more it should do so*' (p. 105; emphasis added). Here we will focus on the last part of this statement, conveying the following comparative principle of confirmation by uncertain evidence:

(H) If coming to believe with certainty that e would confirm h, then, the more probable it becomes that e is true, the more this should confirm h.

Huber considers various Bayesian confirmation measures, provides his own formal analysis of Bayesian confirmation in the 'kilt' case and argues that the difference measure D, the ratio measure R and the odds ratio measure OR all *violate* the allegedly compelling principle (H). He concludes that serious doubts arise on the adequacy of the Bayesian approach and elaborates the point in various ways.

As Huber himself points out, however, his own example can be read in two ways: (i) on one hand, P_y and P_z could be seen as referring to *two alternative possible worlds*, both branching from the state represented by P_x; (ii) on the other hand, P_x, P_y and P_z could be seen as following each other in *a single time sequence*. Importantly, Jeffrey conditionalisation can be indifferently applied if either (i) or (ii) is adopted and, in both cases, it provides one unique value for $P_y(h)$ as well as one unique value for $P_z(h)$.[5] Yet the distinction between the possible worlds and the time sequence interpretation emphasises that principle (H) is equivocal, as it can be taken as reflecting each one of *two very different* adequacy requirements imposed on a candidate measure of confirmation by uncertain evidence c.

If the kilt example is read in terms of possible worlds, then the most natural rendition of (H) given in our terms is:

(H.1) Provided that $P_x(h \mid e) > P_x(h)$, if $P_x(e) < P_y(e) < P_z(e)$, then $c_{x,y}(h) < c_{x,z}(h)$.

In words, this means that the higher the increase from the initial probability of an e confirming h, the higher the confirmatory impact on h will be.

[5] One may doubt that $P_z(h)$ will remain equal when arrived at from P_x versus from P_y, i.e. that

$$P_x(h \mid e)\, P_z(e) + P_x(h \mid \neg e)\, P_z(\neg e) = P_y(h \mid e)\, P_z(e) + P_y(h \mid \neg e)\, P_z(\neg e).$$

This will be so, however, in virtue of a condition known as 'rigidity' (Jeffrey [1965], Chapter 11) or 'invariance' (Jeffrey [2004], p. 52), according to which $P_x(h \mid e) = P_y(h \mid e)$ and $P_x(h \mid \neg e) = P_y(h \mid \neg e)$. It can be proven that rigidity is implied by Jeffrey conditionalisation (indeed, it is *logically equivalent* to it).

If, however, the kilt example is read in terms of a single time sequence (which is Huber's main line in his paper), then principle (H) can also be seen as stating:

(H.2) Provided that $P_x(h \mid e) > P_x(h)$, if $P_x(e) < P_y(e) < P_z(e)$, then $c_{x,y}(h) < c_{y,z}(h)$.

As compared to (H.1), *this* is a completely different claim: it means that *any* subsequent increase (*no matter how small*) in the probability of an *e* confirming *h* will have a greater confirmatory impact on *h* than *any* previous increase (*no matter how large*) in the probability of *e*.

Our claim here is that, while (H.1) is a perfectly safe and sound intuitive constraint on an adequate theory of confirmation by uncertain evidence, (H.2) is utterly implausible (as it will be argued shortly).

As for (H.1), it can be shown that the following holds (see the Appendix for a proof):

Theorem. Any Bayesian confirmation measure $c_{x,y}(h)$ enjoying generalised *P*-incrementality satisfies (H.1).

By contrast, in appropriate cases, all alternative confirmation measures considered here will agree in *violating* (H.2)—as they should, we urge. In fact, it is easy to conceive examples where the increase from $P_y(e)$ to $P_z(e)$ is *so much smaller* (on any plausible standard of comparison) than the increase from $P_x(e)$ to $P_y(e)$ that (H.2) is a highly unappealing principle.

To illustrate, suppose that

$$P_x(e \mid h) = 1$$
$$P_x(h) = 0.05$$
$$P_x(e) = 0.10$$
$$P_y(e) = 0.80$$
$$P_z(e) = 0.81$$

By Jeffrey conditionalisation, it can be computed that

$$P_y(h) = 0.40$$
$$P_z(h) = 0.405$$

Then

$$D_{x,y}(h) = P_y(h) - P_x(h) = 0.35 > 0.005 = P_z(h) - P_y(h) = D_{y,z}(h)$$

Similarly, it can be computed that

$$R_{x,y}(h) = 8 > 1.0125 = R_{y,z}(h)$$
$$OR_{x,y}(h) \approx 12.667 > 1.021 \approx OR_{y,z}(h)$$
$$Z_{x,y}(h) \approx 0.368 > 0.008 \approx Z_{y,z}(h)$$

Thus, all four confirmation measures considered here appropriately violate (H.2) in simple clear-cut cases, i.e. when a subsequent increase in the probability of an e confirming h is unequivocally very small (e.g. 0.80 to 0.81) as compared to a previous increase in the probability of the same e (e.g. 0.10 to 0.80).[6]

We conclude that, contrary to Huber's claim, Bayesian confirmation theory, when properly generalised, actually gets things right when it comes to confirmation by uncertain evidence—i.e. satisfies principle (H.1) and violates (H.2).[7]

Appendix

(H.1) Provided that $P_x(h \mid e) > P_x(h)$, if $P_x(e) < P_y(e) < P_z(e)$, then $c_{x,y}(h) < c_{x,z}(h)$.

Theorem. Any Bayesian confirmation measure $c_{x,y}(h)$ enjoying generalised P-incrementality satisfies (H.1).

Proof. First of all, let's point out that, in what follows, we are positing $P_x(e) > 0$, so that $P_x(h \mid e)$ is defined. Given that, we will prove that, assuming $P_x(h \mid e) > P_x(h)$, $P_x(e) < P_y(e) < P_z(e)$ and $c_{x,y}(h)$ enjoys generalised P-incrementality, the inequality $c_{x,y}(h) < c_{x,z}(h)$ is verified.

6 It is fair to say that this line of argument is partly anticipated, and criticised, by Huber towards the end of his paper ([2005], pp. 111 ff.). Huber's critical point essentially amounts to the remark that, when uncertain evidence is at issue, $c_{x,y}(h)$ crucially depends on P_x even in qualitative terms (confirmation versus disconfirmation). This seems, however, an appropriate feature of Bayesian confirmation by uncertain evidence. Indeed, should it be the case that—for any reason—looking at Stephen actually *decreased A*'s confidence in e down to 0.6 from an initially higher value, we would like to say that this has *disconfirmed h* to some extent. In fact, in such a situation, $c_{x,y}(h)$ would assume a negative value, since by Jeffrey conditionalisation, $P_y(h)$ would itself be lower than $P_x(h)$. (Also see footnote 7.)

7 The lack of an explicit unpacking of statement (H) may not be the only reason why Huber ([2005]) thinks otherwise. From his analysis, it seems that a further reason boils down to his own way of applying Bayesian confirmation under Jeffrey conditionalisation. To illustrate, consider the 'difference' measure of confirmation. In our notation, Huber ([2005], p. 104) seems to have primarily employed the following way of computing degrees of confirmation:

$$D^*_{x,y}(h) = P_y(h \mid e) - P_y(h).$$

This is unfortunate, however, for this quantity does *not* measure the departure of the appropriately updated probability of h from the initial one. In fact, under Jeffrey conditionalisation, it seems obvious that $P_y(h)$, and *not* $P_y(h \mid e)$, represents the degree of belief in h at time y—when the probability of e has shifted to non-extreme values—whereas $P_x(h)$, and *not* $P_y(h)$, represents the initial degree of belief in h. Indeed, if $D^*_{x,y}(h)$ is adopted, not only the implausible principle (H.2), but even the compelling requirement (H.1) itself will be systematically violated. This is bad enough, but it gets worse. For $D^*_{x,y}(h)$ implies that even a *decrease* (!) in the probability of a confirming e will confirm h. In fact, it can be proven that, *for whatever* (non-extreme) *value* of $P_x(e)$ and $P_y(e)$, provided that $P_x(h \mid e) > P_x(h)$, $D^*_{x,y}(h)$ will be higher than the neutrality value 0. In the presence of what we see as a highly plausible alternative way to apply Bayesian confirmation to uncertain evidence, which does not exhibit such undesirable properties, the latter remarks seem to show the inadequacy of $D^*_{x,y}(h)$—not of Bayesian confirmation theory itself.

By the probability calculus, the following equivalence can be derived:

$$P_x(h \mid e) > P_x(h) \Leftrightarrow P_x(h \mid e) > P_x(h \mid \neg e). \tag{1}$$

Since by hypothesis, $P_x(h \mid e) > P_x(h)$, then, by Equation (1), one has $P_x(h \mid e > P_x(h \mid \neg e)$ as well, whence $P_x(h \mid e) - P_x(h \mid \neg e) > 0$. Also, by hypothesis, $P_y(e) - P_x(e) > 0$. So the product of $P_x(h \mid e) - P_x(h \mid \neg e)$ and $P_y(e) - P_x(e)$ will be itself greater than zero. The latter inequality can be algebraically manipulated as follows:

$$\begin{aligned}
&[P_x(h \mid e) - P_x(h \mid \neg e)] \cdot [P_y(e) - P_x(e)] > 0 \Leftrightarrow \\
&P_x(h \mid e) \cdot [P_y(e) - P_x(e)] - P_x(h \mid \neg e) \cdot [P_y(e) - P_x(e)] > 0 \Leftrightarrow \\
&P_x(h \mid e) \cdot [P_y(e) - P_x(e)] + P_x(h \mid \neg e) \cdot [P_y(\neg e) - P_x(\neg e)] > 0 \Leftrightarrow \\
&P_x(h \mid e) P_x(e) + P_x(h \mid \neg e) P_x(\neg e) < P_x(h \mid e) P_y(e) + P_x(h \mid \neg e) P_y(\neg e). \qquad (2)
\end{aligned}$$

By the theorem of total probabilities, Equation (2) can be rewritten as:

$$P_x(h) < P_x(h \mid e) P_y(e) + P_x(h \mid \neg e) P_y(\neg e). \tag{3}$$

Since, by hypothesis, also $P_z(e) - P_y(e) > 0$, an analogous manipulation yields:

$$P_x(h \mid e) P_y(e) + P_x(h \mid \neg e) P_y(\neg e) < P_x(h \mid e) P_z(e) + P_x(h \mid \neg e) P_z(\neg e). \tag{4}$$

And, by Jeffrey conditionalisation, Equations (3) and (4) imply:

$$P_x(h) < P_y(h) < P_z(h). \tag{5}$$

By enjoying generalised P-incrementality, $c_{x,y}(h)$ is by definition a strictly increasing function of the update probability of h. Hence, from Equation (5) it immediately follows that:

$$c_{x,y}(h) < c_{x,z}(h).$$

Acknowledgements

Research supported by PRIN 2005 grant *Le dinamiche della conoscenza nella società dell'informazione*, by a grant from the SMC/Fondazione Cassa di Risparmio di Trento e Rovereto for the CIMeC (University of Trento) research project on *Inductive reasoning*, and by PRIN 2006 grant *Rational Decisions, Strategic Interactions, Complexity and Evolution of Social Systems* (Research unit: Department of Philosophy, University of Trieste). We would like to thank Branden Fitelson and Franz Huber for useful comments on this paper.

Department of Arts and Design
University IUAV of Venice
Dorsoduro 2206
30123 Venice, Italy
crupi@iuav.it

Department of Philosophy
University of Trieste
Androna Campo Marzio 10
34123 Trieste, Italy
festa@units.it

CIMeC
University of Trento
Palazzo Fedrigotti
Corso Bettini 31
38068 Rovereto (TN), Italy
tommaso.mastropasqua@unitn.it

References

Carnap, R. [1950/1962]: *Logical Foundations of Probability*, Chicago: University of Chicago Press.

Christensen, D. [1999]: 'Measuring Confirmation', *Journal of Philosophy*, **96**, pp. 437–61.

Cooke, R. M. [1991]: *Experts in Uncertainty. Opinion and Subjective Probability in Science*, Oxford, UK: Oxford University Press.

Crupi, V., Tentori, K. and Gonzalez, M. [2007]: 'On Bayesian Measures of Evidential Support: Theoretical and Empirical Issues', *Philosophy of Science*, **74**, pp. 229–52.

Earman, J. [1992]: *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge, MA: MIT Press.

Eells, E. [1982]: *Rational Decision and Causality*, Cambridge, UK: Cambridge University Press.

Festa, R. [1999]: 'Bayesian Confirmation', in M. Galavotti and A. Pagnini (*eds*), *Experience, Reality, and Scientific Explanation*, Dordrecht: Kluwer, pp. 55–87.

Fitelson, B. [1999]: 'The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity', *Philosophy of Science*, **66**, pp. S362–78.

Fitelson, B. [2001]: 'A Bayesian Account of Independent Evidence with Applications', *Philosophy of Science*, **68**, pp. S123–40.

Gaifman, H. [1979]: 'Subjective Probability, Natural Predicates and Hempel's Ravens', *Erkenntnis*, **21**, pp. 105–47.

Good, I. J. [1950]: *Probability and the Weighing of Evidence*, London: Griffin.

Good, I. J. [1983]: *Good Thinking*, Minneapolis, MN: University of Minnesota Press.

Horwich, P. [1982]: *Probability and Evidence*, Cambridge, UK: Cambridge University Press.

Huber, F. [2005]: 'Subjective Probabilities as Basis for Scientific Reasoning?', *British Journal for the Philosophy of Science*, **56**, pp. 101–16.

Jeffrey, R. [1965]: *The Logic of Decision*, New York: McGraw-Hill.

Jeffrey, R. [1992]: *Probability and the Art of Judgment*, Cambridge, UK: Cambridge University Press.

Jeffrey, R. [2004]: *Subjective Probability. The Real Thing*, Cambridge, UK: Cambridge University Press.

Joyce, J. [1999]: *The Foundations of Causal Decision Theory*, Cambridge, UK: Cambridge University Press.

Keynes, J. [1921]: *A Treatise on Probability*, London: Macmillan.

Milne, P. [1996]: 'Log[$p(h/eb)/p(h/b)$] is the One True Measure of Confirmation', *Philosophy of Science*, **63**, pp. 21–6.

Mortimer, H. [1988]: *The Logic of Induction*, Paramus, NJ: Prentice Hall.

Nozick, R. [1981]: *Philosophical Explanations*, Oxford, UK: Clarendon Press.

Rescher, N. [1958]: 'A Theory of Evidence', *Philosophy of Science*, **25**, pp. 83–94.

Rips, L. J. [2001]: 'Two Kinds of Reasoning', *Psychological Science*, **12**, pp. 129–34.

van Fraassen, B. [1980]: 'Rational Belief and Probability Kinematics', *Philosophy of Science*, **47**, pp. 165–87.