Bayesian Confirmation:
A Means with No End
Peter Brössel and Franz Huber

ABSTRACT

Any theory of confirmation must answer the following question: what is the purpose of its conception of confirmation for scientific inquiry? In this article, we argue that no Bayesian conception of confirmation can be used for its primary intended purpose, which we take to be making a claim about how worthy of belief various hypotheses are. Then we consider a different use to which Bayesian confirmation might be put, namely, determining the epistemic value of experimental outcomes, and thus to decide which experiments to carry out. Interestingly, Bayesian confirmation theorists rule out that confirmation be used for this purpose. We conclude that Bayesian confirmation is a means with no end.

1 Introduction

Any theory of confirmation must answer the following question: what is the purpose of its conception of confirmation for scientific inquiry, broadly construed? In this article, we argue that no Bayesian conception of confirmation can be used for guiding or prescribing epistemic states, or for guiding or prescribing experimental actions. First we argue that no Bayesian conception of confirmation can be used for its primary intended purpose, which we take to be making a claim about how worthy of belief certain hypotheses are. Then we consider a different use to which confirmation might be put, namely, to determine the epistemic value of experimental outcomes, and thus to decide which experiments to carry out. Interestingly, however, Bayesian confirmation theorists rule out that confirmation be used for this purpose. We conclude that

there is no need to supplement the general Bayesian framework with Bayesian confirmation theory. Bayesian confirmation is a means with no end.[1]

We argue as follows: Bayesian confirmation theory defines confirmation as a relation between evidence, hypothesis, background assumptions, and some probability measure. The probability measure is interpreted as an agent's degree of belief function. Since the agent's degrees of belief are used to determine whether the evidence confirms the hypothesis, confirmation cannot be used to determine the agent's degrees of belief, that is, how worthy of belief the hypothesis is.

Next we consider a popular measure of incremental confirmation and show that it can be used to determine the epistemic value of experimental outcomes. Standard decision theory then allows us to form the expected epistemic value of an experiment, and thus to decide which experiments to carry out. Finally, we look at the notion of confirmation Bayesian confirmation theorists want to explicate. We find that one particular criterion they impose on their explicans, viz. that old evidence should be able to provide confirmation, explicitly rules out that it be used to determine the epistemic value of experimental outcomes. This implies our negative conclusion.

We proceed as follows: In Section 2 we present the tenets of Bayesian confirmation theory. In Section 3 we argue that no Bayesian conception of confirmation can be used to say how worthy of belief certain hypotheses are. In Section 4 we find that there is a Bayesian conception of confirmation that can be used to determine the epistemic value of experimental outcomes. However, we also show that this possible use is ruled out—in some cases explicitly, in others implicitly—by proponents of Bayesian confirmation theory. We conclude in Section 5.

2 Bayesian Confirmation Theory

Bayesian confirmation theory is defined within the general Bayesian framework. The latter consists of the synchronic norm that an agent's degrees of belief should obey the probability calculus and the diachronic norm that,

---

[1] In this article, we focus on showing that no Bayesian conception of confirmation can serve the purpose of guiding or prescribing epistemic states or experimental actions. In other words, we focus on showing that Bayesian confirmation theory, *qua* normative theory, is a means with no end. An anonymous referee for this journal has raised the point that the Bayesian conception of confirmation might serve the purpose of explaining or describing the epistemic relationship between hypotheses and evidence. Even if that were the case, our conclusion—that Bayesian confirmation theory, *qua* normative theory, is a means with no end—would not be affected. Furthermore, in light of work by psychologists such as Kahneman and Tversky (Kahneman and Tversky [1973], Tversky and Kahneman [1983]), we are sceptical that Bayesianism is a descriptively adequate theory. Since it is not clear to us to what extent a descriptively inadequate theory can be explanatory, we are sceptical of the prospects for an explanatory version of Bayesian confirmation theory.

depending on the format of the new information received, the agent's degrees of belief should be updated by strict conditionalization, Jeffrey conditionalization (Jeffrey [1983a]), or Field conditionalization (Field [1978]).

Bayesian confirmation theory adds two notions of confirmation to this general Bayesian framework: absolute confirmation and incremental confirmation. According to both notions, confirmation is a relation between evidence E, hypothesis H, background assumptions B, and some probability measure Pr (and some formal language of sentences or algebra of propositions, which may be ignored for present purposes).

Evidence E is said to absolutely confirm hypothesis H relative to background assumptions B and probability function Pr if and only if the conditional probability of the hypothesis given the evidence and the background assumptions is sufficiently high, $Pr(H \mid E \cap B) > r$, where r is a number not smaller than ½, but smaller than 1. The degree of absolute confirmation of hypothesis H by evidence E relative to background assumptions B is defined as the conditional probability of the hypothesis given the evidence and the background assumptions, $Pr(H \mid E \cap B)$.

Evidence E is said to incrementally confirm hypothesis H relative to background assumptions B and probability function Pr if and only if the conditional probability of the hypothesis given the evidence and the background assumptions is greater than the probability of the hypothesis given merely the background assumptions, $Pr(H \mid E \cap B) > Pr(H \mid B)$. The degree of incremental confirmation of hypothesis H by evidence E relative to background assumptions B is defined as the distance between the conditional probability of the hypothesis given the evidence and the background assumptions, $Pr(H \mid E \cap B)$, and the probability of the hypothesis given merely the background assumptions, $Pr(H \mid B)$. There are several ways to measure distance (Earman [1992], Chapter 5), and there are other quantities the distance between which can be measured (Joyce [2003]). The way distance is measured, and the quantities between which distance is measured, affect the validity of various arguments in Bayesian confirmation theory (Fitelson [1999]; Brössel [2013]). However, these differences do not matter for present purposes.

The probability measure figuring in these definitions of confirmation is typically interpreted as an agent's actual or rational degree of belief function. Subjective Bayesians hold that the probability function is interpreted as the agent's actual degree of belief function, properly idealized. Objective Bayesians hold that the probability function is interpreted as the degree of belief function the agent objectively ought to have (Williamson [2010]), or as the degree of belief function the agent is rationally justified in holding (Maher [2004]).

3 Bayesian Confirmation and Belief

Any theory of confirmation must answer the following question: what is the purpose of its conception of confirmation for scientific inquiry, broadly construed? In this section, we argue that no Bayesian conception of confirmation can be used for its primary intended purpose, which we take to be making a claim about how worthy of belief certain hypotheses are.

Intuitively, we should believe or accept those hypotheses that are well confirmed. For example, suppose the NASA announces that the analysis of data from a moon probe confirms the existence of H_2O. This is generally taken to provide reason to believe or accept the hypothesis that there exists H_2O on the moon. Pre-theoretically, the information that a hypothesis has been confirmed provides reason to believe or accept this hypothesis.

Saying how worthy of belief certain hypotheses are is also the purpose the founding father of confirmation theory intended for confirmation. In his seminal 'Studies in the Logic of Confirmation', Hempel wrote:

> It is now clear that an analysis of confirmation is of fundamental importance also for the study of the central problem of what is customarily called epistemology; this problem may be characterized as the elaboration of 'standards of rational belief' or of criteria of warranted assertibility. (Hempel [1945], p. 7)

We claim that no Bayesian conception of confirmation can be used for this purpose. The reason for this is purely conceptual and consists in the fact that Bayesian confirmation theory presupposes that this question is already answered.

Recall that, according to Bayesian confirmation theory, confirmation is a relation between evidence, hypothesis, background assumptions, and some probability measure. The probability measure is typically interpreted as an agent's (subjective or objective) degree of belief function. This means that confirmation is itself defined in terms of (subjective or objective) degrees of belief. Therefore it cannot be used to say how worthy of belief certain hypotheses are without presupposing that this question is already answered. This is so even if we grant, as we do, that the complete probability function (the complete probability space, if we include the language or algebra), as well as the evidence and background assumptions, are transparent to the agent. The reason is simply that Bayesian confirmation theory interprets the probabilities as the agent's (subjective or objective) degrees of belief.

Let us go through this argument in more detail. In order to say whether evidence E confirms hypothesis H relative to background assumptions B and probability measure Pr, we must specify the probability measure Pr. First consider subjective Bayesians. They interpret the probability measure Pr as the agent's actual degree of belief function, properly idealized. Hence, we must

specify the agent's actual degrees of belief before we can say whether the evidence confirms the hypothesis. Therefore we cannot use the information that the evidence confirms the hypothesis in order to specify the agent's actual degrees of belief.

Next consider objective Bayesians. They interpret the probability measure as the degree of belief function the agent objectively ought to have, or as the degree of belief function she is rationally justified in holding. Hence, we must specify the degrees of belief the agent objectively ought to have, or the degrees of belief she is rationally justified in holding, before we can say whether the evidence confirms the hypothesis. Therefore we cannot use the information that the evidence confirms the hypothesis to say which degrees of belief the agent objectively ought to have, or which degrees of belief she is rationally justified in holding.

Our argument works as long as the probability measure—in terms of which confirmation is defined—is interpreted as the agent's (subjective or objective) degree of belief function, in terms of which confirmation should say how worthy of belief certain hypotheses are. There are two options to escape the negative conclusion of our argument that Bayesian confirmation is a means with no end. The first option is to say that the probability measure in terms of which confirmation is defined—call it the 'confirmation function'—is different from the agent's (subjective or objective) degree of belief function. The second option is to say that confirmation is a means to some end other than saying how worthy of belief certain hypotheses are.

Let us consider the prospects of the first option, according to which the confirmation function is interpreted in terms of something other than the agent's degrees of belief. An attempt along these lines, motivated by the problem of old evidence (see the next section) and championed by Howson and Urbach ([1993], [2006]), is to identify the confirmation function with the agent's counterfactual degree of belief function. The proposal is that

> [...] the support of *H* by *E* is gauged according to the effect which one believes a knowledge of *E would now have* on one's degree of belief in *H*, on the (counterfactual) supposition that one does not yet know *E*. (Howson and Urbach [1993], pp. 404–5, notation adapted; see also Howson and Urbach [2006], pp. 297–301).

The counterfactual strategy as developed by Howson and Urbach ([1993], [2006]) is generally considered unsuccessful as a solution to the problem of old evidence (Williamson [2000], Chapter 10; Hawthorne [2005]). Among other reasons, this is so because it is not clear how to determine the degree of belief an agent would have in *H* given the '(counterfactual) supposition that [she] does not yet know *E*'. That is, it is not clear how to determine the agent's counterfactual degree of belief function from her actual degree of belief

function. However, even if the counterfactual strategy were successful as a solution to the problem of old evidence, it would still not help in escaping the negative conclusion of our argument. The reason is that we need to know the agent's actual degree of belief function before we can determine her counterfactual degree of belief function. Otherwise, one cannot determine 'the effect which one believes a knowledge of *E would now have* on one's degree of belief in *H*'. In order to escape the negative conclusion of our argument, one must interpret the confirmation function as something that is not merely different from, but also sufficiently independent of, the agent's degree of belief function.

One way of doing so is to simply stipulate that there exists a confirmation function that is sufficiently independent of, but nevertheless appropriately linked to, the agent's degree of belief function. We call this strategy of solving philosophical problems by simply postulating solutions to them the *deus ex machina* strategy. The *deus ex machina* strategy is, of course, an unacceptable move in philosophical theorizing. Presumably, no philosopher will ever seriously propose to solve a philosophical problem by it. However, a bad strategy need not be the worst strategy. While there may not be a sharp border separating good strategies from bad ones, other things being equal, strategies are worse the closer they are to the extreme of the *deus ex machina* strategy.

The *deus ex machina* strategy must not be confused with the indispensable move in philosophical theorizing of postulating that a notion of philosophical interest is primitive. For instance, Williamson ([2000]) argues that knowledge should, or at least may, be taken as primitive because attempts to reduce it to notions other than knowledge have failed. He then proceeds to show that other notions of philosophical interest such as assertibility, belief, and evidence can be illuminated by, or even reduced to, knowledge. Another example is Woodward ([2003]), who shows that one can say many illuminating things about direct causal relevance in terms of other causal notions in a circular, but not viciously circular, way. In these two cases, the existence and intelligibility of knowledge and causality (that differs from direct causal relevance), respectively, are taken for granted. While one generally has to argue for the existence and intelligibility of the things one postulates, there are clearer cases such as those of knowledge and causality where this is not so urgent, and less clear cases such as that of the confirmation function where this is urgent. Williamson's work on knowledge and Woodward's work on causality belong to the very best of contemporary philosophy, even though their work bears some resemblance to the *deus ex machina* strategy. It does so by giving up on reductive analysis, and by stipulating that there exist such things as knowledge and causality, and that they have the properties intuition says they have. Therefore it may not be possible to completely do without elements of the *deus ex machina* strategy. However, we clearly should try to work with as few of those elements as possible.

An account that stipulates that there exists a 'degrees of support function' that is sufficiently independent of, and appropriately linked to, the agent's degree of belief function to play the role of the confirmation function is presented in Hawthorne ([2005]). Hawthorne's degrees of support function is supposed to be a logical or quasi-logical conception of probability that represents objective or public likelihoods (Hawthorne [2005], pp. 298ff). While Hawthorne ([2005]) states a link between his degrees-of-support function and the agent's degree of belief function, he does not say why we are entitled to assume that there exists a degrees of support function. Nor does he tell us how we could possibly determine the quasi-logical or logical degree of support of a sentence or proposition. Therefore, Hawthorne violates the criterion of 'ascertainability', which has been suggested by Salmon ([1966]) as a condition of adequacy on interpretations of probability (the criterion is also favourably discussed in Hájek [2012], Section 2): 'This criterion requires that there be some method by which, in principle at least, we can ascertain values of probabilities. It merely expresses the fact that a concept of probability will be useless if it is impossible in principle to find out what the probabilities are' (Salmon [1966], p. 64). Indeed, the only reason Hawthorne provides for the existence of the degrees of support function is that it solves the problem of old evidence. We think that this account bears too much resemblance to the *deus ex machina* strategy.

A similar account is presented in Williamson ([2000], Chapter 10), who stipulates that there exists an 'evidential probability' measure that plays the role of the confirmation function. Having in mind the betting interpretation of an agent's degrees of belief, Williamson explicitly rejects the demand for an operational definition of evidential probability and merely says that it satisfies the axioms of the probability calculus. The only reason he provides for the existence of the evidential probability measure is his dissatisfaction with the notion of an agent's rational degrees of belief. We agree that there is no need for an operational definition of evidential probability. However, we think that the agent's betting behaviour is merely used to measure, not to define, the agent's degrees of belief. In accordance with the above criterion of ascertainability we think that a theory of measurement *is* something highly desirable. Without such a theory, more has to be said about why we are entitled to assume that there exists an evidential probability measure. As with Hawthorne's account of degrees of support, but in stark contrast to Williamson's account of knowledge, we think that Williamson's account of evidential probability bears too much resemblance to the *deus ex machina* strategy.

Hawthorne ([2005]) and Williamson ([2000], Chapter 10) stipulate the existence of a degrees-of-support function and an evidential probability measure, respectively. Taken as the confirmation function, each would be sufficiently independent of, but appropriately linked to, the agent's degree of belief

function in order for confirmation to be a means to the end of saying how worthy of belief certain hypotheses are. However, we do not think these authors provide sufficient reason for us to be entitled to assume the existence of these functions. Their accounts bear too much resemblance to the *deus ex machina* strategy.

We both think that the only notion whose existence and intelligibility may be safely assumed and which is sufficiently independent of, but appropriately linked to, the agent's degree of belief function is the agent's initial or *a priori* degree of belief function. We both also think that this conclusively establishes that there is no need for confirmation as a means to the end of saying how worthy of belief certain hypotheses are. The general Bayesian framework answers this question satisfactorily and without reference to its supplement Bayesian confirmation theory; Bayesian confirmation is a means with no end. (We disagree about what, if anything, one should do about this, though. See (Huber [2005a], [2005b], [2008a], [2008b]; Brössel [unpublished]).)

Therefore let us turn to the prospects of the second option, according to which confirmation is a means to some end other than saying how worthy of belief certain hypotheses are.

4 Confirmation and the Value of Experiments

No Bayesian conception of confirmation can be used to say how worthy of belief certain hypotheses are. However, this still leaves open the possibility that confirmation is put to some other use. In this section, we consider a popular measure of incremental confirmation and show that it can be used to measure the epistemic value of experimental outcomes, and thus to decide which experiments to carry out by forming their expected epistemic values. However, we also show that Bayesian confirmation theorists rule out—in some cases explicitly, in others implicitly—that confirmation be used for this purpose, and so defend our negative conclusion from the previous section.

One might argue that the epistemic value of an experimental outcome in a test of a hypothesis for an agent consists in the difference the experimental outcome would make to the agent's degree of belief in the hypothesis. In the literature, this is referred to as 'the potential further support' (Christensen [1999]) or 'the additional evidence provided by' (Milne [forthcoming]) an experimental outcome for a hypothesis for an agent.

Suppose an agent with actual degree of belief function Pr is testing the hypothesis H by an experiment with possible experimental outcomes E and $\neg E$, against background assumptions B. The difference the experimental outcome E would make to the agent's degree of belief in the hypothesis given the background assumptions, can be measured by the Kemeny–Oppenheim measure of 'factual support', F (Kemeny and Oppenheim [1952]; Good [1983];

Pearl [1988], and Fitelson [2001] endorse F or measures ordinally equivalent to it):

$$F(H, E, B) = \left[\frac{Pr(E \mid H \cap B) - Pr(E \mid \neg H \cap B)}{Pr(E \mid H \cap B) + Pr(E \mid \neg H \cap B)} \right].$$

The difference the experimental outcome $\neg E$ would make to the agent's degree of belief in the hypothesis, given the background assumptions, is measured by

$$F(H, \neg E, B) = \left[\frac{Pr(\neg E \mid H \cap B) - Pr(\neg E \mid \neg H \cap B)}{Pr(\neg E \mid H \cap B) + Pr(\neg E \mid \neg H \cap B)} \right].$$

For an agent with degree of belief function Pr, the expected epistemic value of an experiment $\mathcal{E}$ with possible experimental outcomes E and $\neg E$ for hypothesis H, relative to background assumptions B, can then be defined in the standard decision-theoretic way as

$$\mathcal{F}(H, \mathcal{E}, B) = F(H, E, B)Pr(E \mid B) + F(H, \neg E, B)Pr(\neg E \mid B).$$

There are other measures of incremental confirmation besides F, and it is an interesting question which one is best suited to measuring an agent's epistemic value of an experimental outcome for a hypothesis relative to some background assumptions.

Importantly, though, Bayesian confirmation theorists rule out that any measure of incremental confirmation be used for this purpose. For instance, Christensen ([1999]) wants

> [...] to capture the support an agent's confidence in H *already* receives from E (in contrast to the potential further support that might be gotten from raising $Pr(E)$). (Christensen [1999], p. 449, notation adapted)

He thus explicitly rules out that confirmation be used to measure the epistemic value of experimental outcomes.

Other philosophers follow Christensen in aiming at an explication of 'actualized support' (Christensen [1999], p. 449) as opposed to potential further support. More specifically, many Bayesian confirmation theorists are committed to denying that a measure of confirmation quantifies potential further support, even if they do not deny this explicitly. This can be seen from the fact that they take the following to be a problem for their accounts of incremental confirmation: Recall that evidence E incrementally confirms hypothesis H relative to background assumption B and probability measure Pr just in case $Pr(H \mid E \cap B) > Pr(H \mid B)$. On this account, evidence E that receives probability one relative to background assumptions B cannot confirm or disconfirm any hypothesis H. The reason is that the conditional probability of the hypothesis given the evidence and the background assumptions, $Pr(H \mid E \cap B)$, equals the probability of the hypothesis given merely the

background assumptions, $Pr(H \mid B)$, if the probability of the evidence given the background assumptions equals one, $Pr(E \mid B) = 1$.

Evidence may receive probability one because it is trivial, in the sense that it is highly unspecific or even tautologous. However, on a Bayesian interpretation of probability, evidence may also receive probability one because it is old, in the sense that it has been established beyond doubt. This strips the confirmatory potential from highly non-trivial evidence that is maximally believed to be true. Glymour ([1980]) cites the fact that the pre-existing evidence of Mercury's 43″ arc advance of its perihelion was considered by most physicists to be compelling evidence for Einstein's general theory of relativity as one historical instance of this problem, which since is known as the 'problem of old evidence'.

Numerous attempts have been made to solve the problem of old evidence. They range from claiming that $Pr(E \mid B)$ is never one (Rosenkrantz [1983]) to counterfactually reassigning the probability of old evidence (Howson and Urbach [1993], [2006]), to relocating the source of the problem in logical learning (Garber [1983]; Jeffrey [1983b]; Niiniluoto [1983]).

All these philosophers apparently hold that old evidence is, or should be, able to provide confirmation. More specifically, these philosophers apparently hold that if c is a measure of confirmation, then it is, or should be, possible that there is a piece of evidence, E, and a hypothesis, H, such that E is old given the background assumptions B, $Pr(E \mid B) = 1$, and E confirms H relative to B in the sense of c. However, if c is a measure of potential further support, then c is a measure of incremental confirmation. In this case, it is not possible that there is a piece of evidence E and a hypothesis H such that E is old given background assumptions B, $Pr(E \mid B) = 1$, and E confirms H relative to B in the sense of c. Put differently, for potential further support, the inability of old evidence to raise or lower the probability of some hypothesis is a precondition rather than a problem. A measure of potential further support cannot account for old evidence, nor should it. It cannot and should not do so for the most general reason that old evidence has, by definition, no potential to further support a hypothesis. Old evidence is, by definition, unable to provide additional evidence for any hypothesis.

To sum up: old evidence cannot provide incremental confirmation, and since potential further support is a form of incremental confirmation, old evidence cannot provide potential further support. Therefore all philosophers who take the problem of old evidence seriously cannot want to capture potential further support with their notions of confirmation. This in turn implies that they implicitly rule out that confirmation be used to determine the epistemic value of experimental outcomes, and thus to decide which experiments to carry out.

5 Conclusion

In this article we have argued that no Bayesian conception of confirmation can be used for guiding or prescribing epistemic states or experimental actions. In particular, we have argued that no Bayesian conception of confirmation can be used for its primary intended purpose, which we take to be making a claim about how worthy of belief certain hypotheses are. We have considered a different use confirmation might be put to, namely, to determine the epistemic value of experimental outcomes and so to decide which experiments to carry out. We have seen that Bayesian confirmation theorists rule out that confirmation be used for this purpose, either explicitly or implicitly, by taking the problem of old evidence to be serious. We conclude that Bayesian confirmation is a means with no end.

Peter Brössel
Department of Philosophy
Ruhr-University Bochum
Germany
peter.broessel@rub.de

Franz Huber
Department of Philosophy
University of Toronto
Toronto, Canada
franz.huber@utoronto.ca

References

Brössel, P. [unpublished]: 'Rethinking Bayesian Confirmation Theory – Steps Towards a New Theory of Confirmation'.

Brössel, P. [2013]: 'The Problem of Measure Sensitivity Redux', *Philosophy of Science*, **80**, pp. 378–97.

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

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

Field, H. [1978]: 'A Note on Jeffrey Conditionalization', *Philosophy of Science*, **45**, pp. 361–7.

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]: *Studies in Bayesian Confirmation Theory*, PhD Dissertation, University of Wisconsin-Madison, available at <http://fitelson.org/thesis.pdf>.

Garber, D. [1983]: 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory', in J. Earman (*ed.*), *Testing Scientific Theories*, Minneapolis: University of Minnesota Press, pp. 99–131.

Glymour, C. [1980]: *Theory and Evidence*, Princeton: Princeton University Press.

Good, I. J. [1983]: *Good Thinking: The Foundations of Probability and its Applications*, Minneapolis: University of Minnesota Press.

Hájek, A. [2012]: 'Interpretations of Probability', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy* <http://plato.stanford.edu/archives/win2012/entries/probability-interpret/>.

Hawthorne, J. [2005]: 'Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions', *Mind*, **114**, pp. 277–320.

Hempel, C. G. [1945]: 'Studies in the Logic of Confirmation (I)', *Mind*, **54**, pp. 1–26.

Howson, C. and Urbach, P. [1993]: *Scientific Reasoning: The Bayesian Approach*, La Salle: Open Court.

Howson, C. and Urbach, P. [2006]: *Scientific Reasoning: The Bayesian Approach*, 3rd edition, La Salle: Open Court.

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

Huber, F. [2005b]: 'What Is the Point of Confirmation?', *Philosophy of Science*, **72**, pp. 1146–59.

Huber, F. [2008a]: 'Assessing Theories, Bayes Style', *Synthese*, **161**, pp. 89–118.

Huber, F. [2008b]: 'Milne's Argument for the Log-Ratio', *Philosophy of Science*, **75**, pp. 413–20.

Jeffrey, R. C. [1983a]: *The Logic of Decision*, Chicago: University of Chicago Press.

Jeffrey, R. C. [1983b]: 'Bayesianism with a Human Face', in J. Earman (*ed.*), *Testing Scientific Theories*, Minneapolis: University of Minnesota Press, pp. 133–56.

Joyce, J. M. [2003]: 'Bayes' Theorem', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/fall2008/entries/bayes-theorem/>.

Kahneman, D. and Tversky, A. [1973]: 'On the Psychology of Prediction', *Psychological Review*, **80**, pp. 237–51.

Kemeny, J. and Oppenheim, P. [1952]: 'Degree of Factual Support', *Philosophy of Science*, **19**, pp. 307–24.

Maher, P. [2004]: 'Probability Captures the Logic of Scientific Confirmation', in C. Hitchcock (*ed.*), *Contemporary Debates in the Philosophy of Science*, Oxford: Blackwell, pp. 69–93.

Milne, P. [unpublished]: 'Measuring Confirmation', available at <www.fitelson.org/few/milne.pdf>.

Niiniluoto, I. [1983]: 'Novel Facts and Bayesianism', *British Journal for the Philosophy of Science*, **34**, pp. 375–9.

Pearl, J. [1988]: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, San Francisco: Morgan Kaufman.

Salmon, W. [1966]: *The Foundations of Scientific Inference*, Pittsburgh: University of Pittsburgh Press.

Rosenkrantz, R. D. [1983]: 'Why Glymour is a Bayesian', in J. Earman (*ed.*), *Testing Scientific Theories*, Minneapolis: University of Minnesota Press, pp. 69–97.

Tversky, A. and Kahneman, D. [1983]: 'Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement', *Psychological Review*, **90**, pp. 292–315.

Williamson, J. [2010]: *In Defence of Objective Bayesianism*, Oxford: Oxford University Press.

Williamson, T. [2000]: *Knowledge and Its Limits*, Oxford: Oxford University Press.

Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, Oxford: Oxford University Press.