13

INTRODUCTION

Science is founded on the supposition that the evidence makes some theories reasonable to accept, and others unreasonable. Some hypotheses, that is, are supported by the evidence, but not all. Moreover, we have a pretty good idea of which hypotheses those are: for example, evolution, plate tectonics, the theory of relativity. And we have a pretty good idea of how to subject new hypotheses to evidential test. We have a broad competence, that is, with the scientific method, and reasonable agreement between scientists about how to deploy it and interpret its results.

On the other hand, there are famous philosophical arguments to the effect that observational evidence will never support any particular beliefs about the future, or any theoretical hypothesis about the unobserved more generally, and that inductive or ampliative inference is not rational (Hume, 1748). So, if there is no reasonable inference from evidence to hypothesis, how can it be that science is successful and there is a broad agreement on its methods? To explain this, we need to develop some understanding of evidential support that either goes beyond or evades Hume's arguments.

Famously, Popper (1959) agreed with Hume and concluded that there could be no reasonable grounds to accept a theory on the basis of evidence. Rather, we could at best *reject* theories that turned out to be inconsistent with the evidence (where inconsistency is the ordinary logical notion). This doctrine of *falsificationism* is still a dominant trope in scientists' own conception of the scientific method; unfortunately it fits very poorly with scientific practice—often it poorly fits the practice even of those who espouse it. For most scientists do think that past evidence and the success of science in predicting new evidence provides excellent reason to believe a scientific theory, not just good reason to refrain from rejecting it, as Popper would have it. So our delicate position continues: science is manifestly successful, vindicating scientific inference; but we have philosophical reason to suspect that there are no good theories of rational inference that would validate scientific reasoning.

The views on credence and updating we have explored in the previous two parts suggest a way out of this Humean impasse. As we saw, there is a sense in which our credences represent our confidence in particular propositions; and there is a good sense in which updating our credences on the basis of new evidence represents our subsequent confidence in those propositions. What if, by updating on new evidence, we become more confident in a proposition than we were previously? In that case, it seems, the evidence we have received has supported the proposition in question. Hume's arguments—if they succeed—would show that there is no rationally required inductive response to evidence. The *Bayesian* alternative, where 'Bayesian' is intended to capture broadly probabilist views on credence and updating, is to say that the rational updating of coherent credences on receipt of new evidence is rational. But how one responds to evidence depends on one's prior conditional credences, which may vary from person to person, and those priors are not in any sense rationally mandated. We shall look more closely at these kinds of Humean arguments in Part IV, where views that dispute Hume's contention that there is no a priori support for inductive reasoning will be discussed.

Confirmation

The Bayesian account of confirmation focuses on the inferences of individual scientists when confronted with evidence. As suggested above, the proposal is that, if C is the scientist's prior credence, and C^+ their posterior credence, then an experience provides evidence that *confirms* a hypothesis h, for that scientist, if $C^+(h) > C(h)$. If the posterior credence is assumed to obey conditionalisation on the evidence proposition e, as discussed in Part II, this is equivalent to $C(h|e) > C(h)$. This evidence 'boosts' the scientist's credence in the hypothesis.

It might initially be suspected that, while this might explain the inferences of an individual scientist, the variety of permissible degrees of belief that the subjective probability framework permits will undermine any attempt to explain general community-wide patterns of scientific inference—the patterns which have collectively been called the 'scientific method'. But, in fact, from what seem like relatively minimal and uncontroversial assumptions on prior credences, many of the accepted canons of scientific inference can be given a reconstruction within the Bayesian framework. Insofar as those constraints on priors are widely followed, there is therefore a reason for most to follow the scientific method, which is thereby given a justification. Of course, the justification here is that the scientific method encapsulates in simple heuristics the inferences which could have been directly established by conditionalising on the evidence, so it is hardly independent reason for anyone to adopt prior credences which obey the required constraints. (If (i) agents are not perfectly introspective and cannot tell whether or not they satisfy the constraints, and (ii) there is advantage to being in line with community standards with respect to confirmation, as it seems there might well be in actual scientific communities, then acting so as to follow the rules, and thereby ensure that one acts as if one's credences are of the appropriate sort, might be prudentially justified.) The justification on offer, then, is dependent on the fact that—as things actually stand—we do tend to have credences with the right kind of structure.

This has seemed like grounds for imposing additional constraints on prior credence. So Williamson, in arguing for a kind of objective Bayesianism, in which there is a unique rational initial credence distribution which 'measures something like the intrinsic plausibility of hypotheses prior to investigation' (the kind of view we will consider in Part IV), writes that 'successful [subjective] Bayesian treatments of specific epistemological problems ... assume that subjects have "reasonable" prior distributions. We judge a prior distribution reasonable if it complies with our intuitions about the intrinsic plausibility of hypothesis. This is the same sort of vagueness as infects the present approach, if slightly better hidden' (Williamson, 2000: 211–12). The successes of Bayesianism, in effect, are claimed to be by-products of the prior and have little to do with the minimal requirement of coherence. (A similar observation is made by Glymour, Chapter 15, p. 256, where he argues that an *ad hoc* assignment of priors can be constructed that will ratify any particular canons of scientific inference. But we

didn't want to show that the scientific method was *consistently implementable* in the Bayesian framework—we wanted to show that it followed from all (or most) reasonable credences that these inferences were good ones.)

Yet it is arguably not reasonable to expect something as strong as a unique prior assignment of intrinsic plausibility, as it seems manifestly plausible that people can reasonably disagree over the prior plausibility of a hypothesis. What matters more, surely, is posterior agreement, or at least agreement on what confirms what. The examples developed by Howson and Urbach in Chapter 14 seem to show that the goal of explaining the substantive agreement on confirmation amongst scientists can be met by subjective Bayesians who appeal to the *de facto* structural similarity between scientists' priors. But this goal is equally satisfied by the theories of evidential probability discussed in Chapter 17, particularly the strand that takes the existence of strong constraints on prior conditional probabilities as a starting point. And note, moreover, that whether the agreement on the prior conditional probabilities is given by a rationality constraint or just what is institutionally or congenitally true of us, the same patterns of explanation of the maxims in the examples below will obtain.

Before developing some of those examples, I will mention an argument some Bayesians offer that the subjectivity of the priors doesn't matter. This argument points to the existence of 'convergence of opinion' theorems (discussed by Glymour in Chapter 15; also mentioned in Chapter 1), to the effect that two Bayesian agents who start off with regular priors (more generally, who assign credences of 1 and 0 to the same propositions as each other), and who successively update by conditionalisation on the same sequence of evidence propositions will, in the limit with probability 1, come to have posterior credences that are arbitrarily close to one another, no matter where they started out. The differences in the priors are swamped by the agreement on the evidence. It can also be shown that, if the evidence sequences are relatively simple, and the hypotheses under consideration similarly simple (perhaps just being hypotheses about the total evidence sequence), these agents will converge not just on each other, but will both converge on the true hypothesis in the limit, with probability 1. So in this case it seems that the subjectivity of the priors doesn't matter, because these two agents will come to agree in their final judgements about which hypotheses are most plausible in light of the evidence. But we never have infinite sequences of data, so there seems no guarantee from these theorems that actual people will come to agree. It is an equally false idealisation to assume that two agents might have the same evidence—even ideal agents will differ at least in their *de se* (self-locating) evidence. Whatever explains our actual general agreement, it can't be anything like what is involved in the conditions for the convergence of opinion theorems to hold, and so appeal to those theorems seems unlikely to provide a justification for the methodology actually employed in scientific inference. Finally, a point that Glymour makes is worth noting: the relevant convergence theorem

> does not tell us that in the limit any rational Bayesian will assign probability 1 to the true hypothesis and probability 0 to the rest; it only tells us that rational Bayesians are certain that he will. It may reassure those who are already Bayesians, but it is hardly grounds for conversion.
>
> (Chapter 15: 255)

Even if it is thereby rational to expect convergence and that differences in the priors don't matter, that doesn't yet show that these rational expectations will be fulfilled.

Successes of the Bayesian approach to confirmation

A 'success' for Bayesian confirmation theory is a case where, if the pre-theoretical notion of evidential support is analysed in Bayesian terms, the apparently rational responses to evidence that are licensed by the informal notion can be motivated and reconstructed in the Bayesian framework. There can generally be two types of such cases: *assumption-free* cases where the reconstruction succeeds regardless of the prior credences involved, and *assumption-laden* cases where the reconstruction requires some assumptions about the priors. The latter cases will also permit the Bayesian to construct circumstances where agents have coherent prior credences that diverge from those needed to reconstruct the target inference. Such cases will provide another direction of support for the Bayesian analysis of confirmation if it can be shown that, in these cases where the Bayesian predicts it is not necessarily rational to respond to evidence in a certain familiar way, our intuitive judgements also make the same prediction. (This latter kind of case is arguably what has happened with the Bayesian response to Hempel's paradox of confirmation, where Good (1961) gave a case that persuaded many that a prima facie plausible attempt to capture in a maxim some of the intuitive features of confirmation by instances had an unexpected counterexample, the existence of which is revealed by a Bayesian approach.)

The approach of Howson and Urbach in Chapter 14 is to offer Bayesian reconstructions and explanations of a number of widely accepted doctrines about evidence in science. For example, they show how Bayesians can explain the following familiar maxims of scientific interference:

Surprising evidence. Surprising evidence e supports a hypothesis h which predicts it—that is, assigns it a high *likelihood* $C(e|h)$—over those which do not predict it. Moreover, the more surprising e is, other things being equal, the higher the posterior credence in h will be (pp. 226–230).

Paradox of Confirmation. It seems that generalisations are confirmed by their instances ('Nicod's condition'), and logically equivalent propositions are each confirmed by the same evidence. But it follows from this that the observation of a white shoe, confirming as it does the generalisation 'All non-black things are non-ravens', should thereby confirm the logically equivalent generalisation 'all ravens are black'. Yet it does not seem to. On pages 204–6, Howson and Urbach argue, following Hosiasson-Lindenbaum (1940), that in the Bayesian framework the paradoxical conclusion does not follow (for Nicod's condition is not generally satisfied). Even under those circumstances when Nicod's condition is satisfied, and the observation of a white shoe does confirm the hypothesis that all ravens are black, it will typically (i.e., given ordinary background assumptions) only raise the posterior probability by a negligible amount, and much less than the boost offered by the observation of a black raven. They take the comparatively weak evidential support offered by the observation of a white shoe to explain our ordinary judgements about the problem cases; we fail to distinguish negligible support from no support at all.

Duhem's problem. Duhem and Quine argued that a hypothesis can be brought into contact with evidence only with the aid of auxiliary claims (so, for example, a theory of the motion of the planets can only make predictions about what we observe through the telescope with the aid of an auxiliary theory

of how telescopes work). *Duhem's problem* is to rationalise our rejection (in most cases) of the hypothesis rather than the auxiliary; on pages 230–236, Howson and Urbach, following the work of Dorling (1979), provide a Bayesian rationale.

***Ad hoc* hypotheses.** It is often maintained that *ad hoc* theories constructed to fit the data are less believable than independently motivated theories; relatedly, that predicting a piece of evidence provides more support for a theory than being constructed so as to accommodate a piece of evidence. The discussion on pages 238–245 summarises the Bayesian rationale for these practices.

Some of these cases, like the case of surprising evidence, are assumption-free; the treatment of *ad hoc* theories, on the other hand, is assumption-laden. Even so, the fact that Bayesianism can provide a plausible theoretical framework in which all of these various canons of scientific inference can be argued for is already a mark in its favour. Certainly there is no current alternative general theory of what confirmation consists in that can explain them, which is (at least by the Bayesian's own lights) some support for the Bayesian analysis of confirmation. Rather than go through the details of these examples, which are clearly laid out in Chapter 14 and dealt with more fully in the items cited in Further Reading below, I will illustrate the kind of reconstruction on offer from Bayesianism with reference to another example: the evidential value of diverse evidence (see also Chapter 14, pp. 244–5).

Imagine we are considering a hypothesis h. The likelihood of the evidence is the conditional credences assigned to evidence, conditional on the hypotheses $C(e|h)$. The *Bayes factor* of $\neg h$ against h is defined as this ratio:

$$\beta(\neg h : h) = \frac{C(e|\neg h)}{C(e|h)}.$$

It can be shown that we can rewrite Bayes' theorem (Probability Primer, theorem 13) in terms of the Bayes factor and the prior probability of h (the prior probability of e going unmentioned):[1]

$$C(h|e) = \frac{C(h)}{C(h) + \beta(\neg h : h)C(\neg h)}. \qquad \text{(Useful)}$$

This form is useful, at least, because many have thought that the likelihoods are less 'subjective' than the other credences (it is commonly thought that the likelihood of evidence conditional on a theory is something that is determined by the content of the theory, and what evidence it predicts); this form shows where the updated credence ends up in terms of the prior of the hypothesis and these less subjective likelihoods. (Of course, since the likelihoods replace an expression referring to the prior credence in the evidence, which is supposed to be subjective, perhaps this claim about the objectivity of likelihoods is not sustainable.)

It appears to be a methodological rule that, other things being equal, the more diverse the sources of evidence for one's theory, the more strongly confirmed that theory is. This can be captured in this maxim: *A theory which makes predictions in a number of disparate and seemingly unconnected areas is more confirmed by that evidence than* *is a theory which is confirmed by predictions only about a narrow and circumscribed range*. The Bayesian insight is that diverse evidence is not internally correlated (Earman, 1992: Sect. 3.5). If the evidence is diverse, it consists of at least two propositions, e_1 and e_2, such that truth of one is not positively relevant to the truth of the other if the hypothesis in question is false. So e_1 and e_2 are diverse relative to h iff the likelihood $C(e_1 \wedge e_2|\neg h)$ is low, or at least if it is not greater than the product of the individual credences $C(e_1|\neg h)C(e_2|\neg h)$. If, for example, the hypothesis is that all swans are white, then swans collected from different countries would, if white, provide better evidence for the hypothesis than swans collected from the same pond, as we know that if one swan on a pond is white, it is much more likely to be related to other swans in its pond, and those are more likely therefore to be white.

If the hypothesis h predicts both e_1 and e_2, then the likelihood $C(e_1 \wedge e_2|h)$ is high (close to one). So the Bayes' factor is approximately equal to its numerator, $C(e_1 \wedge e_2|\neg h)$. Substituting this into (Useful), we get

$$C(h|e) \approx \frac{C(h)}{C(h) + C(e_1 \wedge e_2|\neg h)C(\neg h)}$$
$$\geq \frac{C(h)}{C(h) + (C(e_1|\neg h)C(e_2|\neg h))C(\neg h)}.$$

The second line follows from the diversity of the evidence. Thus far, we are only drawing out consequences of the probability calculus. In practical cases, we are interested in evidence that is diverse in such a way as to ensure that, if the hypothesis is false, at least one of the pieces of evidence is unlikely. It would be surprising, given that not all swans are white, if arbitrary swans taken from diverse locations were all likely to be white. There is an additional assumption on the priors in this case: that at least one of $C(e_1|\neg h)$ and $C(e_2|\neg h)$ is low, and perhaps both are low. Given that, $C(h) + (C(e_1|\neg h)C(e_2|\neg h))C(\neg h)$ may be approximately $C(h)$ and therefore $C(h|e)$ may be close to 1. If h was not antecedently plausible, diverse evidence of this form has strongly confirmed it. (A similar rationale exists for the practice of *random sampling*, if the idea is to ensure pieces of evidence that are uncorrelated with one another.)

Problems for Bayesianism

The picture for Bayesianism is bright, but far from unalloyed. Glymour in Chapter 15 proffers a number of difficulties for Bayesianism, additional to those already raised above. The first part of his chapter primarily concerns the assignment of degrees of belief and the implementation of respect for simplicity in priors; but the final part of his chapter concerns the Bayesian account of evidence.

Glymour raises a number of problems for that account. One interesting puzzle (at p. 271) for the understanding of science more generally is that, if we can (even in a rough and ready way) separate propositions into those directly about observation, and those which are theoretical, then we can construct a 'rival' h' to any theory h, which is just the conjunction of the observational consequences of h. These two theories are equally confirmed by the evidence, since they predict the same appearances; but since $h \rightarrow h'$, the posterior credence in h' cannot be lower than the posterior in h, and hence h' is at least as credible. But this doesn't seem right: this kind of rival is never seriously considered in science, and the theoretical virtues of simplicity, explanatory power, etc.,

which h' lacks and h may possess, do play a role in scientific inference. The Bayesian, at least in this respect, seems to flout standard scientific practice.

But it is another worry—the problem of *old evidence*—which has been a major stumbling block for Bayesianism. This is the simple observation that, while 'scientists commonly argue for their theories from evidence known long before the theories were introduced' (p. 262), and that even in these cases there can be striking confirmation of a new theory because it explains some well-known anomalous piece of evidence, the Bayesian framework doesn't seem to permit this. For if e is old evidence, then one's credence function at the time the new theory h is proposed, C_t, has come from conditionalising on e at some stage. But then $C_t(h|e) = C_t(h)$, so that the posterior credence in h on e is equal to the prior; h is not confirmed.

Glymour's own diagnosis of the problem is that Bayesianism is (at best) a theory of *learning*, of how to represent your credence at a time and how to update your credence when you learn something new. Sometimes our evidence is something we learn; but sometimes, as in these cases, we have already learned our evidence. The relation of evidential support between evidence and hypotheses seems as though it should be indifferent to whether we know the evidence before or after considering the theory; by elevating the accidental fact that sometimes we learn our evidence into an account of evidence, Bayesianism is unable to accommodate this.

One response is to go on the offensive: if the evidence really is old, then all the scientists already believe it, and have already updated their credences in the hypothesis in question. If there seems to be additional confirmation in these cases, it is because those scientists have imperfect access to their own credences; or are incoherent, since (despite already knowing e) they are more confident in h after noticing e, even though no new evidence has arrived. These Bayesians say: if confirmation doesn't have a close connection with learning, that only undermines its importance—for the main aim of scientific inference isn't to see what confirms what for its own sake, but to discover what we should believe, and what we believe is constrained by the theory of subjective probability.

There are a variety of Bayesian responses which are more concessive than this. Two responses share the idea that, while the old evidence doesn't support the hypothesis given the agent's current credences, there are other credences which are appropriate in some way and on which the evidence does confirm the hypothesis, because it is learned. One response appeals to counterfactual credences—what our credences would have been, had we not learned e. The other appeals to past credences—what our credences were, before we learned e. And the claim is that h is confirmed by e against the background of these counterfactual or historical credences. Both responses seem to run into difficulty from the same kind of case. Take a very old piece of evidence, well-known and very familiar, which, in conjunction with some auxiliary claim a that we learned much more recently, turn out to support a new hypothesis. But if we return to our credence before we learned e, we abandon not only e but also a; and it needn't be that h is confirmed by e without a. (Consider: we've known that the earth was round since classical times, but that doesn't mean that the way in which that fact supports hypotheses must involve only auxiliaries also known since classical times.) If, furthermore, a was only learned because e was learned (e is a striking fact about some subject matter, which prompted further investigation of that area and related areas, which led to the discovery of a), then our counterfactual credence again needn't confirm h because a is not known. (The nearest possibility in which we hadn't learned e is also one where we hadn't learned a either.)

An alternative approach (favourably discussed by Glymour at pp. 265–6) is to say that, while *e* is already certain, what is learned is that *h* entails *e*. Garber (1983) argues that this is best considered a failure of logical omniscience, and that it is possible that Bayesian agents who fail to be certain of all logical truths can, when they learn such a truth, come to be aware of new confirmation relations. But liberalising knowledge of logic threatens to undermine the whole Bayesian approach (since, for example, those who are not logically omniscient can be Dutch booked), quite apart from the difficulties involved in trying to construct a theory that permits logically equivalent propositions to have different credences.

Old evidence remains a live problem for Bayesianism. But other problems have been raised too. The problem of *new theories* arises when a well-confirmed theory is undermined by the postulation of a new hypothesis, even though no new evidence has been received and no credential update has occurred. Again, this kind of case could be solved by appeal to (logical or introspective) ignorance, since a new theory in this sense should still be one which is a proposition over which the credence function is defined, and its novelty must derive from its being unrecognised. A more radical proposal is that this might be a case in which an new algebra of propositions must be considered; it is fair to say that changing the algebra of outcomes is a non-conditionalising update which Bayesians have no account of, though it is equally unclear whether anyone else does either.

Various alternative theories of confirmation also pose challenges for the Bayesian. The main one is the classical theory of statistical testing, developing from the work of Fisher (and ultimately representing the application of Popperian ideas to probabilistic hypotheses); a good recent account of how this project bears on confirmation theory is offered by Mayo (1996). The computational epistemology defended by Kelly (2003) does away with the idea that evidence justifies hypotheses, replacing it with the notion of a reliable procedure for accommodating evidence to converge on the truth. The goodness of such a procedure is evaluated on the externalist criterion of whether or not it leads to the truth. This is a radical alternative to internalist conceptions of rationality, as captured by Bayesianism's emphasis on explaining scientific inferences from within the framework of one's own priors. Once there are other options on the table, it may be that there are better justifications for the scientific method than just that our priors happen to have the right form. Finally, Norton (forthcoming) provides a useful compendium of problems for the Bayesian approach to confirmation.

Measures of confirmation

So far the focus has been on the qualitative relation of confirmation, whether *e* confirms *h*. But there is a follow-up question: *how much* does *e* confirm *h*? Does *e* confirm *h more than* it confirms *h*′? Many of the maxims of scientific inference involve implicit appeal to notions like this. In the case of diverse evidence, what really seems to matter is that diverse evidence confirms a theory more than narrow evidence does. Another example is provided by Howson and Urbach's attempt to explain away our intuition that a non-black non-raven does not confirm the hypothesis that all ravens are black. They claim that, while there is confirmation in some cases like Hempel's,

> Once it is recognised that confirmation is a matter of degree, the conclusion is no longer so counter-intuitive, because it is compatible with [a non-black non-raven] confirming 'All ravens are black', but to a minuscule and negligible degree.
>
> (Chapter 14, page 228)

Hopefully, any plausible measure of confirmation will agree with their assessment of the degree of confirmation in these cases, but it is important to recognise that an adequate Bayesian response to the paradox will require some account of degree of confirmation that will yield this result.

Moreover, if degree of confirmation could be accounted for in the Bayesian framework, it would provide another answer to the subjectivity of the priors—for it could perhaps be argued that, if *e* supports *h* much more than it supports *h*′, there is a reason to believe *h* independently of the prior probability of these hypotheses.

Eells and Fitelson (Chapter 16) address both the issue of how to define the *degree of confirmation* and the role that notion plays in Bayesian reconstructions of the scientific method. They proceed by noting three symmetry constraints on degree of confirmation, intuitively motivated by consideration of the examples they describe in Sect. 3 of their chapter. Our intuitions about strength of evidence are, as seen in many cases, quite as strong as our intuitions about qualitative confirmation. So the failure or holding of these symmetry constraints provides natural and intuitive support for those measures of confirmation that agree with those judgements. Eells and Fitelson argue that, of the natural measures that have been proposed in the literature, only the difference measure *d* (that *e* confirms *h* to the degree that the posterior exceeds the prior) and the log-likelihood measure *l* (that *e* confirms *h* to a degree measured by the log of the ratio of the likelihood of *e* on *h* to the likelihood of *e* on $\neg h$) satisfy all three of their symmetry constraints. Insofar as these symmetry constraints are motivated by the counterexamples, the field is narrowed to those two measures. (Note that their positive argument for the symmetry condition HS is not strictly needed, as the other measures are ruled out by their satisfaction of other intuitively violated symmetry constraints.)

These observations are not unimportant. As Fitelson (1999) points out, several Bayesian treatments of issues of interest, including the problem of diverse evidence, seem to depend on properties of the measure of confirmation, and the success of these treatments then depends on the measure of confirmation having the right form.

Those who favour other measures can either try to respond to the cases on offer (and those on offer in Fitelson 1999), or argue that since there are intuitive considerations in favour of their preferred measures too, there is no univocal notion of degree of support here to be explicated. This approach has bold and reasonable varieties. The bold variety would be to argue that the intuitive notion of degree of support is inconsistent, supporting many constraints that cannot be jointly sustained. (It would be analogous to the way in which, it might be argued, natural language seems to support the T-schema and the existence of names for arbitrary sentences of the language, giving rise to the inconsistencies involved in the Liar paradox.) The more reasonable variety is *pluralist* about degree of confirmation; there are many overlapping but distinct notions of evidence and support in play, all of which have a home in the Bayesian framework, and each of which has a role to play in coming to understand the complexities of pre-theoretical intuitions about confirmation. Again, the debate over which, if any, measure of confirmation is 'the' right one is far from settled.

Further reading

The idea that Bayesian confirmation theory provides a concessive way out of Hume's problem of induction is clearly presented in Howson (2000). A recent discussion of the connection between confirmation theory and induction, with reference to Howson's book, is Strevens (2004). Accounts of how Bayesian confirmation theory bears on the so-called 'new riddle' of induction (Goodman 1954) include those of Sober (1994) and Fitelson (2008).

The further reading in Chapter 1 contains details of convergence of opinion theorems.

Perhaps in 1980 it was fair of Glymour to complain that 'There is very little Bayesian literature about the hotchpotch of claims and notions that are usually canonised as scientific method' (Chapter 15, p. 256), but the picture has changed dramatically in the interim. General discussions of Bayesian confirmation theory are provided by Earman (1992), Fitelson (2001), Horwich (1982), Jeffrey (2004: Ch. 2), Talbott (2008: Sect. 4), Rosenkrantz (1977), Sober (2002), and Strevens (2006). A broader discussion of the consequences of Bayesian principles in epistemology is Bovens and Hartmann (2003).

The paradox of confirmation arises obviously in logical theories of confirmation, like that of Hempel (1945). (A useful recent discussion of such theories, aimed at rehabilitating something in the neighbourhood of what Hempel wanted, is Huber (2008a)). The Bayesian counterexample to the principles involved was offered by Good (1967). Vranas (2004) provides a thorough account of the standard Bayesian solutions, like that of Horwich (1982: 54–63), and notes an apparent gap. Fitelson (2006) gives an accessible recent overview, containing (as does the article by Vranas) extensive references to the literature, and offering at the end a patch for the gap identified by Vranas. Maher (1999) provides an account of the paradox from a Carnapian/logical probability perspective.

The orthodox Bayesian solution to Duhem's problem is by Dorling (1979); a recent account, under some conditions, is Bovens and Hartmann (2003: Sect. 4.5). A revised version of the Bayesian account is proposed by Strevens (2001), which is disputed by Fitelson and Waterman (2005).

Regarding *ad hoc* hypotheses, the example of Neptune is discussed by Jeffrey (2004: Sect. 2.3). Two disagreeing but nevertheless orthodox Bayesian treatments of prediction versus accommodation are by Horwich (1982: 108–18) and Maher (1988). A clear recent discussion of the methodological precept that Bayesians are supposedly capturing is White (2003).

Horwich (1982: 118–22) offers a more general account of the value of diverse evidence, predicated upon the ability of diverse evidence to rule out competitors to a hypothesis that agree with the hypothesis on some narrower body of evidence. The standard Bayesian solution is criticised by Wayne (1995); replies are offered by Fitelson (1996) and Steel (1996).

The Bayesian difficulty with confirmation of theories, as opposed to summaries of their observational consequences, is turned into an argument for anti-realism about scientific theories by van Fraassen (1980), who appeals to the same fact that theories are conjunctions of their observational content and theoretical content. Van Fraassen (1989: Ch. 7) offers a Dutch book argument against inference to the best explanation, and other realist patterns of scientific inference. There is a huge literature on this kind of argument for anti-realism about science; some pertinent discussions are offered by Douven (1999), Lipton (1991) and Milne (2003).

The historical or counterfactual credences approach to the problem of old evidence is discussed by Jeffrey (2004: Sect. 2.5), drawing on a series of papers by Wagner that culminate in Wagner (2001). Lange (1999) argues that a 'rational reconstruction' of the ur-credence function can permit a Bayesian account of old evidence. The relaxation of the logical omniscience condition advocated by Garber is also advocated by Niiniluoto (1983) and further discussed by Jeffrey (1983); a critical evaluation is Earman (1992: Ch. 5). A useful account of different threads in the discussion of old evidence is in Joyce (1999). The problem of new hypotheses was raised by Chihara (1987); a Bayesian response is provided by Otte (1994). The relation between old evidence and new theories is discussed by Zynda (1995).

The computational learning theory alternative to Bayesianism is put forward and used to argue that Bayesianism can't capture all of scientific methodology by Kelly and Glymour (2004).

Eells' treatment of the grue paradox where d plays some role is to be found in Eells (1982). Fitelson (2001) offers considerations in favour of the log-likelihood measure l. Eells and Fitelson have extensive references to the literature on measures of confirmation in their chapter. The pluralist approach is defended in a paper by Joyce (2003); his 2008 contribution (Sect. 3) contains further hints of his pluralism. Milne (1996) defends an explicitly non-pluralist account of degree of support, though Huber (2008b) shows that a minor variation of Milne's argument supports an alternative measure of confirmation. Further useful discussions are by Christensen (1999) and Eells and Fitelson (2000). The psychology of judgements of degree of confirmation, looking at how a real subject's judgements conform to the measures proposed, is explored by Tentori *et al.* (2007).

Note

1 Proof:

$$C(h|e) = \frac{C(e|h)C(h)}{C(e)} = \frac{C(h)}{\left(\frac{C(e)}{C(e|h)}\right)} = \frac{C(h)}{\left(\frac{C(e|h)C(h)+C(e|\neg h)C(\neg h)}{C(e|h)}\right)}$$
$$= \frac{C(h)}{C(h)+\left(\frac{C(e|\neg h)}{C(e|h)}C(\neg h)\right)} = \frac{C(h)}{C(h)+\beta(\neg h:h)C(\neg h)}$$

Bibliography

Bovens, Luc and Hartmann, Stephan (2003), *Bayesian Epistemology*. Oxford: Oxford University Press.

Chihara, Charles (1987), 'Some Problems for Bayesian Confirmation Theory'. *British Journal for the Philosophy of Science*, vol. 38: pp. 551–560.

Christensen, David (1999), 'Measuring Confirmation'. *Journal of Philosophy*, vol. 96: pp. 437–461.

Dorling, Jon (1979), 'Bayesian Personalism, the Methodology of Scientific Research Programmes, and Duhem's Problem'. *Studies in History and Philosophy of Science*, vol. 10: pp. 177–187.

Douven, Igor (1999), 'Inference to the Best Explanation Made Coherent'. *Philosophy of Science*, vol. 66 (Proceedings): pp. S424–S435.

Earman, John (1992), *Bayes or Bust?*. Cambridge, MA: MIT Press.

Eells, Ellery (1982), *Rational Decision and Causality*. Cambridge: Cambridge University Press.

Eells, Ellery and Fitelson, Branden (2000), 'Measuring Confirmation and Evidence'. *Journal of Philosophy*, vol. 97: pp. 663–672.
Fitelson, Branden (1996), 'Wayne, Horwich, and Evidential Diversity'. *Philosophy of Science*, vol. 63: pp. 652–660.
—— (1999), 'The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity'. *Philosophy of Science*, vol. 66 (Proceedings): pp. S362–S378.
—— (2001), *Studies in Bayesian Confirmation Theory*. PhD thesis, University of Wisconsin, Madison, URL http://fitelson.org/thesis.pdf.
—— (2006), 'The Paradox of Confirmation'. *Philosophy Compass*, vol. 1: pp. 95–113.
—— (2008), 'Goodman's "New Riddle" '. *Journal of Philosophical Logic*, vol. 37: pp. 613–643.
Fitelson, Branden and Waterman, Andrew (2005), 'Bayesian Confirmation and Auxiliary Hypotheses Revisited: A Reply to Strevens'. *British Journal for the Philosophy of Science*, vol. 56: pp. 293–302.
Garber, Daniel (1983), 'Old Evidence and Logical Omniscience in Bayesian Confirmation Theory'. In John Earman (ed.), *Testing Scientific Theories*, Minneapolis: University of Minnesota Press.
Good, I. J. (1961), 'The Paradox of Confirmation (III)'. *British Journal for the Philosophy of Science*, vol. 11: pp. 63–64.
—— (1967), 'The White Shoe is a Red Herring'. *British Journal for the Philosophy of Science*, vol. 17: p. 322.
Goodman, Nelson (1954), *Fact, Fiction and Forecast*. Cambridge, MA: Harvard University Press.
Hempel, Carl (1945), 'Studies in the Logic of Confirmation'. *Mind*, vol. 54: pp. 1–26, 97–121.
Horwich, Paul (1982), *Probability and Evidence*. Cambridge: Cambridge University Press.
Hosiasson-Lindenbaum, Janina (1940), 'On Confirmation'. *Journal of Symbolic Logic*, vol. 5: pp. 133–148.
Howson, Colin (2000), *Hume's Problem: Induction and the Justification of Belief*. Oxford: Oxford University Press.
Huber, Franz (2008a), 'Hempel's Logic of Confirmation'. *Philosophical Studies*, vol. 139: pp. 181–189.
—— (2008b), 'Milne's Argument for the Log-Ratio Measure'. *Philosophy of Science*, vol. 75: pp. 413–420.
Hume, David (1748 [1999]), *An Enquiry Concerning Human Understanding*. Oxford Philosophical Texts, Oxford: Oxford University Press. Edited by Tom L. Beauchamp.
Jeffrey, Richard C. (1983), 'Bayesianism with a Human Face'. In John Earman (ed.), *Testing Scientific Theories*, vol. X of *Minnesota Studies in the Philosophy of Science*, Minneapolis: University of Minnesota Press, pp. 133–156.
—— (2004), *Subjective Probability (The Real Thing)*. Cambridge: Cambridge University Press.
Joyce, James M. (1999), *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
—— (2003), 'On the Plurality of Probabilist Measures of Evidential Relevance', URL http://www.uni-konstanz.de/ppm/kirchberg/Joyce_1.pdf. Unpublished manuscript, presented at the Bayesian Epistemology workshop at the 26th International Wittgenstein Symposium, Kirchberg, Austria, 5 August 2003.
—— (2008), 'Bayes' Theorem'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Fall 2008 ed.
Kelly, Kevin T. (2003), 'The Logic of Success'. In Peter Clark and Katherine Hawley (eds.), *Philosophy of Science Today*, Oxford: Oxford University Press, pp. 11–38.
Kelly, Kevin T. and Glymour, Clark (2004), 'Why Probability does not Capture the Logic of Scientific Justification'. In Christopher Hitchcock (ed.), *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell, pp. 94–114.
Lange, Marc (1999), 'Calibration and the Epistemological Role of Bayesian Conditionalization'. *Journal of Philosophy*, vol. 96: pp. 294–324.
Lipton, Peter (1991), *Inference to the Best Explanation*. London: Routledge.

Maher, Patrick (1988), 'Prediction, Accommodation, and the Logic of Discovery'. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, vol. 1: pp. 273–285.
—— (1999), 'Inductive Logic and the Ravens Paradox'. *Philosophy of Science*, vol. 66: pp. 50–70.
Mayo, Deborah (1996), *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
Milne, Peter (1996), 'log[$P(h|eb)/P(h|b)$] is the One True Measure of Confirmation'. *Philosophy of Science*, vol. 63: pp. 21–26.
—— (2003), 'Bayesianism v. Scientific Realism'. *Analysis*, vol. 63: pp. 281–288.
Niiniluoto, I. (1983), 'Novel Facts and Bayesianism'. *British Journal for the Philosophy of Science*, vol. 34: pp. 375–379.
Norton, John D. (forthcoming), 'Challenges to Bayesian Confirmation Theory'. In Prasanta Bandyopadhyay and Malcolm Forster (eds.), *Handbook of the Philosophy of Science, vol. 7: Philosophy of Statistics*, Amsterdam: Elsevier. URL: http://www.pitt.edu/~jdnorton/papers/Challenges.pdf
Otte, Richard (1994), 'A Solution to a Problem for Bayesian Confirmation Theory'. *British Journal for the Philosophy of Science*, vol. 45: pp. 764–769.
Popper, Karl (1959), *The Logic of Scientific Discovery*. London: Hutchinson.
Rosenkrantz, R. D. (1977), *Inference, Method and Decision: Towards a Bayesian Philosophy of Science*. Dordrecht: D. Reidel.
Sober, Elliott (1994), 'No Model, No Inference: A Bayesian Primer on the Grue Problem'. In Douglas Stalker (ed.), *GRUE! The new riddle of induction*, Chicago: Open Court, pp. 225–240.
—— (2002), 'Bayesianism: its Scope and Limits'. In Richard Swinburne (ed.), *Bayes' Theorem*, Oxford: Oxford University Press.
Steel, Daniel (1996), 'Bayesianism and the Value of Diverse Evidence'. *Philosophy of Science*, vol. 63: pp. 666–674.
Strevens, Michael (2001), 'The Bayesian Treatment of Auxiliary Hypotheses'. *British Journal for the Philosophy of Science*, vol. 52: pp. 515–537.
—— (2004), 'Bayesian Confirmation Theory: Inductive Logic or Mere Inductive Framework?' *Synthese*, vol. 141: pp. 365–379.
—— (2006), 'Notes on Bayesian Confirmation Theory', http://www.nyu.edu/classes/strevens/BCT/BCT.pdf. Unpublished manuscript.
Talbott, William (2008), 'Bayesian Epistemology'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Fall 2008 ed., http://plato.stanford.edu/archives/fall2008/entries/epistemology-bayesian/
Tentori, Katya, Crupi, Vincenzo, Bonini, Nicolao and Osherson, Daniel (2007), 'Comparison of Confirmation Measures'. *Cognition*, vol. 103: pp. 107–119.
van Fraassen, Bas C. (1980), *The Scientific Image*. Oxford: Oxford University Press.
—— (1989), *Laws and Symmetry*. Oxford: Oxford University Press.
Vranas, Peter (2004), 'Hempel's Raven Paradox: a Lacuna in the Standard Bayesian Solution'. *British Journal for the Philosophy of Science*, vol. 55: pp. 545–560.
Wagner, Carl (2001), 'Old Evidence and New Explanation III'. *Philosophy of Science*, vol. 68 (Proceedings): pp. S165–S175.
Wayne, Andrew (1995), 'Bayesianism and Diverse Evidence'. *Philosophy of Science*, vol. 62: pp. 111–121.
White, Roger (2003), 'The Epistemic Advantage of Prediction over Accomodation'. *Mind*, vol. 112: pp. 653–683.
Williamson, Timothy (2000), *Knowledge and its Limits*. Oxford: Oxford University Press.
Zynda, Lyle (1995), 'Old Evidence and New Theories'. *Philosophical Studies*, vol. 77: pp. 67–95.