

The Bayesian Treatment of Auxiliary Hypotheses

Michael Strevens

British Journal for the Philosophy of Science, 52, 515–537, 2001

Copyright British Society for the Philosophy of Science

Abstract

This paper examines the standard Bayesian solution to the Quine-Duhem problem, the problem of distributing blame between a theory and its auxiliary hypotheses in the aftermath of a failed prediction. The standard solution, I argue, begs the question against those who claim that the problem has no solution. I then provide an alternative Bayesian solution that is not question-begging and that turns out to have some interesting and desirable properties not possessed by the standard solution. This solution opens the way to a satisfying treatment of a problem concerning ad hoc auxiliary hypotheses.

1. Introduction

Pierre Duhem (1954) long ago convinced philosophers of science that a scientific hypothesis is never refuted by the evidence. Rather, what is refuted is, even in the simplest cases, the conjunction of the hypothesis with one or more auxiliary hypotheses that concern conditions under which the evidence was produced and observed. Duhem's insight—that a piece of evidence can never refute a single scientific hypothesis, only a conjunction of hypotheses—is sometimes transformed into a far more radical claim, that a piece of evidence can never *test* a single hypothesis, only a conjunction of hypotheses. The principal assumption made in the transition from Duhem's thesis to the stronger thesis, which I will call the Quine-Duhem thesis, is that when a conjunction of a central hypothesis and one or more auxiliary hypotheses is refuted, there is no principled way to distribute blame among the conjuncts, and thus that it is impossible to say to what degree the refutation disconfirms the central hypothesis.

Those who hope to show that the Quine-Duhem thesis is mistaken, then, must answer the question: when a conjunction is refuted, what determines the degree to which each conjunct is disconfirmed? This I will call the Quine-Duhem problem. The Quine-Duhem thesis, then, is the claim that the Quine-Duhem problem is insoluble. The Quine-Duhem problem is interesting not just for its own sake, but because a comprehensive solution to the problem would amount to a general method for introducing auxiliary hypotheses

into confirmation theory. Since, as Duhem observed, auxiliary hypotheses are indispensable in bringing evidence to bear on scientific theories, this is something very much to be desired.

One of the many attractive features of Bayesianism is that it offers a rule of confirmation that easily takes account of auxiliary hypotheses, and so promises to solve the Quine-Duhem problem and thereby to refute the Quine-Duhem thesis. This promise, however, has not yet been fully realized, or so I will argue in section 2 of this paper. I then present what I consider to be the correct Bayesian solution to the Quine-Duhem problem (sections 3 and 4). Later sections discuss the usefulness of the solution for confirmation theory more generally (section 5), and apply it to a particular problem raised by ad hoc auxiliary hypotheses (section 6).

2. The Standard Bayesian Solution to the Quine-Duhem Problem

2.1 The Bayesian Theory of Confirmation

According to the Bayesian theory of confirmation, confirmation and disconfirmation are to be understood as consequences of the norms of rationality governing subjective probabilities. A piece of evidence confirms a hypothesis if, upon learning of the evidence, one ought to increase one's subjective probability for the hypothesis. The degree of confirmation is captured by the degree to which the probability ought to increase. Similarly, a hypothesis

is disconfirmed by a piece of evidence if one's subjective probability for the hypothesis ought to decrease upon learning of the evidence.

At the heart of Bayesian confirmation theory is Bayes' rule of conditionalization, which dictates how one's subjective probabilities ought to change upon learning of any particular piece of evidence.¹ According to Bayes' rule, when one learns that a piece of evidence e has been observed, one's new probability for a hypothesis h , which I will write $P^+(h)$, ought to be set equal to one's old conditional probability for h given e , written $P(h|e)$. That is, one ought to ensure that

$$P^+(h) = P(h|e)$$

What makes Bayes' rule interesting is a theorem of the probability calculus (Bayes' theorem) according to which

$$P(h|e) = \frac{P(e|h)}{P(e)}P(h)$$

Using the theorem, the new probability for h can be calculated from the old probability for h together with the old probabilities for e and e given h .

This last probability, $P(e|h)$, can often be set according to rules generally agreed to be objectively compelling. If h entails e , then $P(e|h)$ must be set to one. If h entails $\neg e$, then $P(e|h)$ must be set to zero. Finally, if h explicitly

1. There exist Bayesian personalists who deny that conditionalization deserves a place in Bayesianism; they hold that the probability calculus, and in particular Bayes' theorem, are the sole constraints on rational belief. But because these Bayesians put no constraints on the change of belief over time, their brand of Bayesianism cannot be regarded as a foundation for confirmation theory.

assigns some physical probability to e , $P(e|h)$ ought, in most circumstances, to be set equal to that physical probability.² For more on the foundations and mechanics of Bayesianism, I refer the reader to Howson and Urbach (1993); Earman (1992).

An important piece of terminology: one's subjective probabilities immediately before learning e are called one's *prior probabilities* by Bayesians.³ One's subjective probabilities immediately after learning e are called one's *posterior probabilities*. Throughout the paper, prior probabilities are written as $P(\cdot)$, posterior probabilities as $P^+(\cdot)$.

2.2 Bayesianism and the Quine-Duhem Problem

Suppose that some piece of evidence e refutes the conjunction of a central hypothesis h and an auxiliary hypothesis a , that is, that ha entails $\neg e$. Because $P(e|ha)$ ought, according to the probability calculus, to be zero, by Bayes' rule one's subjective probability for ha ought to go to zero on learning that e obtains. But how should e affect the probabilities of h and a individually?

2. The kind of rule that enjoins setting subjective probabilities equal to physical probabilities in the manner indicated is what I have called elsewhere a principle of probability coordination (Strevens 1999). Several slightly different probability coordination rules have been described in the literature. Perhaps the best known is David Lewis's "Principal Principle" (Lewis 1980).

3. These are the *relative* priors—relative, that is, to the time as which e is learned. Bayesians also use the expression "prior probability" to refer to what might be called the initial priors, the subjective probabilities that one has before one learns anything about the world.

To provide a method for “apportioning the blame” between h and a is the Quine-Duhem problem. The Quine-Duhem thesis states that the problem cannot be solved, which, within the context of a probabilistic confirmation theory, amounts to the claim that there is no justifiable method for fixing on any particular posterior probability for h or a in the light of the falsification of ha .

Bayesian confirmation theory appears to show, in an utterly straightforward way, that the Quine-Duhem thesis is false. For Bayes’ rule provides a method for calculating the posterior probability of anything given anything else. In particular, given an e that falsifies ha , the effect of e on h and a is encapsulated in the conditional probabilities $P(h|e)$ and $P(a|e)$. Thus $P^+(h)$ may be obtained from Bayes’ rule as follows:

$$\begin{aligned}
 P^+(h) &= P(h|e) \\
 &= P(ha|e) + P(h\neg a|e) \\
 &= P(h\neg a|e) \quad (\text{since } ha \text{ entails } \neg e) \\
 &= \frac{P(e|h\neg a)P(h\neg a)}{P(e)} \\
 &= \frac{P(e|h\neg a)P(h\neg a)}{P(e|h\neg a)P(h\neg a) + P(e|\neg ha)P(\neg ha) + P(e|\neg h\neg a)P(\neg h\neg a)}
 \end{aligned}$$

The formula for $P^+(a)$ simply switches the h s and the a s.

On what I call the standard Bayesian solution to the Quine-Duhem problem, this formula alone is held to be sufficient to solve the problem. That

$P^+(h)$ and $P^+(a)$ are, given the Bayesian framework, to be calculated as shown above is uncontroversial. The claim that this calculation solves the Quine-Duhem problem, however, requires two additional assumptions:

1. The Bayesian method is in general justified, and
2. Given that the Bayesian method is justified, the posterior probabilities $P^+(h)$ and $P^+(a)$ are accurate measures of the impact of the falsification of ha on h and a .

In what follows, I question both assumptions. The first assumption has, of course, been much discussed, both with respect to and independently of the role of Bayes' rule in solving the Quine-Duhem problem. My aim in raising the question is not to refute Bayesian confirmation theory, for I provide a Bayesian solution to the Quine-Duhem problem myself, but to refine the Bayesian response to the critique. As far as I know, the second assumption has not previously been questioned.

2.3 Questioning the First Assumption

A number of writers have complained that the probabilities that determine the value of $P^+(h)$ in the formula presented in the last section are “insufficiently objective” (Earman 1992; Worrall 1994). What is meant is that there is no objectively rational procedure, nor even an intersubjectively agreed on procedure, for setting such probabilities as $P(e|h-a)$ and $P(h-a)$, hence that

there can be no objectively rational procedure, nor even an intersubjectively agreed on procedure, for assessing the impact on h and a of ha 's falsification.

In response to this sort of objection, Bayesians have traditionally provided three different kinds of replies:

1. Personalism: A more objective confirmation theory is not possible.

One must simply learn to accept that confirmation is relative to subjective probability judgments. Howson and Urbach (1993) often, but not always, take this line.

2. Prior Objectivism: There are in many cases objectively justified rules

for assigning values to all the necessary prior probabilities, that is, to all the subjective probabilities that play a role in Bayesian confirmation theory but are not set by Bayesian conditionalization itself (Rosenkrantz 1981; Jaynes 1983).

3. Posterior Objectivism: There are no objectively justified rules for as-

signing values to the prior probabilities. But whatever the values of the priors, Bayesian conditionalization can be shown to produce posterior probabilities with certain objectively desirable properties. A very simple example: if a hypothesis h entails a piece of evidence e , then Bayesian conditionalization on e always raises the probability of h , thus h counts as confirmed by e , as intuition would seem to demand.

Another well-known and considerably more far-reaching argument in this vein is the “washing out” argument, according to which, after a

sufficient amount of evidence comes in, it is almost certain that all Bayesian conditionalizers will have beliefs as to the identity of the true hypothesis that are (a) similar, and (b) correct (Edwards et al. 1963). See Earman (1992) for a critical examination of the assumptions required for the washing out argument to go through.

Prior and posterior objectivism can be combined to considerable effect: by putting even mild constraints on the priors, far stronger results about the posteriors can be obtained than would otherwise be possible. Even washing out arguments assume that the truth is assigned a non-zero prior probability.

The Bayesian literature on the Quine-Duhem thesis does not attempt any of the above replies. Instead, Dorling (1979) has shown that, for certain historically interesting choices of the values of the probabilities in the above formula, the posterior probabilities $P^+(h)$ and $P^+(a)$ act in objectively desirable ways.⁴ Howson and Urbach (1993, §7.h) add further historical arguments of the same sort. From such demonstrations, it can be concluded that, if the historical estimates of the prior probabilities are correct, then Bayesian confirmation theory correctly models the historical probability shifts. But this, of course, does not tell us much about the Quine-Duhem problem, which demands that an objective basis for the probability shifts be provided. Indeed,

4. Of course, Dorling simply relies on philosophical intuition about the nature of confirmation to justify the claim of objective desirability for the behavior. It is assumed that we already know, at least in broad outline, how a confirmation theory ought to behave. I will not question this assumption here.

one does not even have much of an explanation for the intersubjectivity of the probability shifts, since consensus concerning the prior probabilities is assumed rather than accounted for.

In later sections, I hope to do better than this. I will present a kind of posterior objectivist argument: that on the correct Bayesian understanding of the Quine-Duhem problem, Bayesian conditionalization provides a way to assess the impact on a hypothesis h of the falsification of ha that behaves in certain objectively desirable ways, whatever the values of the prior probabilities. Before I do this, however, I will argue that the standard Bayesian solution to the Quine-Duhem problem is incorrect (section 2.4). I then show, in section 2.5, that given the standard, incorrect Bayesian solution to the Quine-Duhem problem, no posterior objectivist argument of the sort I intend to provide would be possible.

2.4 Questioning the Second Assumption

I now argue that the second assumption required by the standard Bayesian solution to the Quine-Duhem problem is false, that is, that it is not always true that the posterior probabilities $P^+(h)$ and $P^+(a)$ are accurate measures of the impact of the falsification of ha on h and a .

My case rests on the following observation: the posterior probability $P^+(h)$ calculated in the standard solution takes into account not only the effect of the falsification of ha , but also ways in which e impacts on h that are independent of the falsification of ha . To spell out the steps leading to

this conclusion: e contains not only the information that ha is false, but other information as well (assuming that $\neg(ha)$ does not entail e). Some of this other information may impact the probability for h for reasons that have nothing to do with the fact that ha is false. Thus the posterior probability $P^+(h)$ will reflect not only the information that ha is false, but also any other relevant information that may be contained in e .

To represent this fact in symbols, I will write:

$$P^+(h) = P(h) + \delta_{qd} + \delta_c$$

where δ_{qd} is the change in the probability of h due to the falsification of ha alone, and δ_c is the change in the probability of h due to factors other than the falsification of ha . (Note that δ_c also represents any effect on the probability of h due to the interaction of the falsification of ha with other information in e .) Please note that this formulation is not intended as a claim about the mathematical structure of the probability change due to e . In particular:

1. I do not claim that the quantity δ_c has some unitary meaning of its own, or that it can be determined independently of the determination of δ_{qd} . Rather, the formula provided above ought to be regarded as the definition of δ_c .
2. I do not claim that the quantities δ_{qd} and δ_c are in some natural sense additive.

3. I do not claim that the change in the probability of h can be determined by two separate acts of conditionalization, one determining δ_{qd} and one determining δ_c , which can then be added to obtain $P^+(h)$.

All I intend is the claim that δ_{qd} is well defined, that is, that there is some method—concerning the mathematical properties of which I make no assumptions at this stage—for determining the part of the probability change for h that is due to the falsification of ha alone.

It is important to see that δ_c need not equal zero, and that it may exert a significant effect on $P^+(h)$, possibly quite different from the effect produced by δ_{qd} . To see that this is so, consider the following scenario. Newstein, a brilliant but controversial scientist, has asserted both that h is true and that e will be observed. You do not know Newstein's reasons for either assertion, but if one of her claims turns out to be correct, that will greatly increase your confidence that Newstein is putting her brilliance to good use and thus your confidence that the other claim will also turn out to be correct. Because of your knowledge of Newstein's predictions, then, your $P(h|e)$ will be higher than it would be otherwise.⁵ This is all as it should be in the Bayesian

5. Here it is important that $P(h|e)$ tacitly includes the effect of all relevant background knowledge, so that e is either reconceived to include both whatever was observed *and* all relevant background knowledge or that $P(h|e)$ is understood to be shorthand for $P(h|ek)$, where k is the relevant background knowledge. In any case, it is generally agreed both (a) that Bayes' rule must incorporate the effect of background knowledge, so that the correct form of the rule is, if e includes only what was observed, $P^+(h) = P(h|ek)$, and (b) that it is desirable in most circumstances to suppress the reference to k , as I do throughout this paper, if only for the purposes of readability.

universe: Bayesians not only are able to, but ought to, take such factors as Newstein's opinions into account. The increased $P(h|e)$ will mean that your $P^+(h)$ is also higher than it would be otherwise. Because the increase is not due to e 's refutation of ha , it is a part of δ_c , not of δ_{qd} . Thus δ_c is not always zero, and it may be positive or—as in the case where Newstein claims that h is false—negative.

The standard Bayesian solution to the Quine-Duhem problem consists in showing that, when ha is refuted, the probability of h can change. This is supposed to show that the refutation of ha does have a quantifiable effect on the probability of h . But the demonstration is incomplete. The standard solution shows that $\delta_{qd} + \delta_c$ is well defined, but it does not follow that δ_{qd} is well defined, unless one assumes that at least one, and therefore both, of the summands are well defined. To make this last assumption is, however, to beg the question, that is, to assume that there exists some method for distinguishing the impact on h of the falsification of ha .

A Bayesian might reply that, in the scenarios sketched by Dorling and others, there are no Newstein effects. More generally, the probabilities have been chosen in such a way that δ_c is zero, so that the entire probability change can be attributed to δ_{qd} . But how is one to evaluate this claim? It is perhaps possible to rule out a Newstein as the source of the probability change (though not entirely if historically informed probabilities are chosen, as historical figures did not ignore one another's opinions). There are, how-

ever, a multitude of ways in which e can bear on h other than by refuting ha .⁶ I do not see how this issue can be discussed clearly and without begging the question in the absence of a generally applicable technique for distinguishing the contributions of δ_{qd} and δ_c to a probability change.

Let me stress that I am arguing a point of principle. I do not believe that Newstein effects are common. Indeed, in section 5 I will argue that, in many important cases, δ_c is close to zero, and this conclusion will play an indispensable role in my treatment of ad hoc hypotheses in section 6. But to simply assert that δ_c is negligible, rather than founding it in some principle, is to rely on “intuition” to discern the relative contributions of δ_{qd} and δ_c to the change in the probability of h . A proponent of the Quine-Duhem thesis will reject the intuition. What is required is an *argument* that δ_c is negligible. Such an argument must, of course, assume a formula for δ_c , and so assume that δ_c , and therefore δ_{qd} , are well defined. In short, the argument must assume that there is a solution to the Quine-Duhem problem, and so cannot itself play a role in such a solution.

The challenge to Bayesianism, then, is to find in the Bayesian apparatus

6. Dorling and Howson and Urbach concede this point. Howson and Urbach (1993) discuss an unpublished paper in which Dorling confronts the fact that there are choices of probabilities on which e falsifies ha but increases the probability of h . Howson and Urbach endorse Dorling’s explanation of this phenomenon, which is the same as mine: “the refuting evidence normally contains a good deal more information than is required merely to disprove $[ha]$ and . . . this extra information may be confirmatory” (Howson and Urbach 1993, 142, paraphrasing Dorling). Apparently neither Howson and Urbach nor Dorling anticipate the problem that I have posed.

a method for calculating not, as the standard solution does, $\delta_{qd} + \delta_c$, but rather δ_{qd} in isolation. This turns out to be quite easy; it is done for the case where ha is falsified in section 3. In section 4, I provide a Bayesian answer to the more general question: how does e impact on h by way of confirming or disconfirming (but not falsifying) ha ?

There is one other reply that a Bayesian might make to my objection. They might hold that, for the purposes of confirmation, it is sufficient that the value of $\delta_{qd} + \delta_c$ be well defined, and so that $P^+(h)$ be well defined. Why ask for more than a posterior probability for h ? This view will, I expect, be appealing to more personalist Bayesians.

In response to this kind of personalist I offer not a refutation but an observation: the personalist is not solving the Quine-Duhem problem, but rather rejecting it. Presumably no proponent of the Quine-Duhem thesis would deny that a piece of evidence that refutes ha can affect h and a in well-defined ways *if* the evidence contains other information that bears directly on h or a . What they deny is that the refutation of ha alone can have a well-defined, rationally motivated effect on h and a . To answer this doubt, one must address the question of the impact of the falsification of ha alone on h , which is, in the Bayesian context, to say that one must give a definition, in Bayesian terms, of δ_{qd} .

The personalist rejecter of the Quine-Duhem problem simply expresses a lack of interest in distinguishing the effect of the falsification of ha from the effect of the opinions of a Newstein. Rather than try to pique this

interest directly, let me point out that, in sections 3 and 4, I provide a quite straightforward and utterly Bayesian way of making the distinction. The personalist can have the solution if they want it. I leave it to them to decide if it is worth the effort.

2.5 Objectivity Again

Let me conclude this section by briefly assessing the significance of Newstein effects for the attempt to defend the standard Bayesian solution to the Quine-Duhem problem from worries about objectivity.

Consider what I have called prior objectivism, that is, the attempt to fix objectively justified values for the prior probabilities that play a role in any particular probability shift. In order to provide an objective basis for the standard Bayesian solution to the Quine-Duhem problem, the prior objectivist must supply rules for fixing every prior that plays a role in determining the posterior probability for a given hypothesis h . The possibility of Newstein effects shows that, among other things, such a rule would have to supply all probabilities necessary for assessing the significance of Newstein's predictions, for example, a likelihood such as the probability that Newstein is right concerning the truth of h given that Newstein is genuinely brilliant and not merely crazy. This is a heavy demand to make; certainly, no account of objectivist priors in the literature seems adequate to satisfy the demand.

Now consider what I have called posterior objectivism, that is, the attempt to show that, regardless of the priors, posterior probabilities will be-

have in certain objectively desirable ways. In order to provide an objective basis for the standard Bayesian solution to the Quine-Duhem problem, the posterior objectivist must show that when ha is falsified, the probability of h changes in a way that has some objectively desirable properties. But this is impossible, simply because, without constraints on the priors, it is impossible to show that the probability of h will change in any particular way. The reason is that a Newstein effect may be engineered so as to change the probability of h in whatever way you like. More generally, if the priors are unconstrained, they may be arranged so as to generate just about any value for δ_c that you like. The fact, by itself, that ha has been falsified, has no ramifications whatsoever for the posterior probability of h . The probability of h may even increase (see note 6).

For these reasons, I am inclined to think that Bayesianism's critics are right in questioning the objectivity of the standard solution to the Quine-Duhem problem. The difficulties posed for both prior objectivism and for posterior objectivism would be ameliorated if there were grounds for ruling out, in some cases at least, the existence of strong Newstein effects, or even better, if there were grounds for showing that in some cases δ_c as a whole was negligible. To do so, however, requires some means for breaking up the probability change for h into different parts. My complaint about the standard Bayesian solution to the Quine-Duhem problem is that it provides no such means. Thus the objections to the standard solution's first and second assumptions are intertwined.

3. Refuting Evidence

Suppose that ha entails $\neg e$, but that e is observed. My aim in this section is to answer the question: how does the observation of e impact on the probability of h in virtue of *and only in virtue of* falsifying ha ?

The quantity I am after is what, in the last section, I called δ_{qd} . It is easier, however, to pursue the closely related quantity $P(h) + \delta_{qd}$. I call this quantity the *partial posterior probability* of h ; it is what the posterior probability of h would be if e only impacted on h by falsifying ha . Because e may affect h in other ways as well, the actual posterior for h may not equal the partial posterior. In such cases, the partial posterior may not be anyone's subjective probability for h . I will write it $Q(h)$ to remind the reader that it is not necessarily a subjective probability.

What is the importance of $Q(h)$, then? Simply that a method for calculating an exact value for $Q(h)$ will solve the Quine-Duhem problem, since an exact value for $Q(h)$ determines an exact value for δ_{qd} . I should make it quite clear that I claim no other significance for the partial posterior. In particular, I do not claim that Bayesians should abandon posterior probabilities and care only about partial posteriors. I cannot emphasize strongly enough that the proposals made in this paper do not revise the mechanics of Bayesian confirmation theory one bit; rather, they first (sections 3 and 4) provide a technique for the interpretation of certain Bayesian probability shifts, and second (sections 5 and 6) provide a technique for demonstrating some of the

properties of those same probability shifts. The shifts themselves are to be determined in an entirely orthodox Bayesian manner.

There is an obvious Bayesian technique for figuring the effect that e has on the probability of h solely in virtue of falsifying ha : take the probability of h conditional on ha being false, that is, conditional on $\neg(ha)$. Then

$$\begin{aligned} Q(h) &= P(h|\neg(ha)) \\ &= \frac{P(\neg(ha)|h)}{P(\neg(ha))}P(h) \end{aligned}$$

I will be interested in finding a revealing expression for $P(\neg(ha)|h)/P(\neg(ha))$, the multiplier for $P(h)$. First, the numerator, $P(\neg(ha)|h)$: if h is true then $\neg(ha)$ can only obtain if $\neg a$. So $P(\neg(ha)|h) = P(\neg a|h)$, which can in turn be written $1 - P(a|h)$. Second, the denominator:

$$\begin{aligned} P(\neg(ha)) &= 1 - P(ha) \\ &= 1 - P(a|h)P(h) \end{aligned}$$

Putting it together:

$$Q(h) = \frac{1 - P(a|h)}{1 - P(a|h)P(h)}P(h)$$

The end product is far simpler than the formula for $P^+(h)$. But how plausible is it?

In the remainder of this section I will offer what is analogous to a posterior objectivist justification for my formula for the partial posterior $Q(h)$, which

is to say, I will show that, regardless of the relevant priors, $Q(h)$ has some properties that are objectively desirable in any measure of δ_{qd} .

First, the multiplier for $P(h)$ is always less than one, so δ_{qd} —the change in the probability of h due to the refutation of ha —is always negative. This confirms the intuition that the falsification of ha is bad for h and a . (Exceptions occur at extreme values: for example, when $P(h)$ is already one, δ_{qd} is zero, as one would expect.)

Second, δ_{qd} depends on the prior probabilities of h and a in the right sort of way. The higher $P(h)$, the less h is blamed when ha is refuted, that is, the smaller the decrease in the probability of h due to the refutation, or in other words, the smaller the absolute magnitude of δ_{qd} . Conversely, the higher $P(a)$, the more h will tend to be blamed when ha is refuted, that is, the greater the decrease in the probability of h .⁷ These properties of the

7. Here I assume that, as $P(a)$ increases, $P(a|h)$ will tend to increase as well (although not necessarily proportionally). This assumption can fail in particular circumstances, but when h and a are a theory and an auxiliary hypothesis, these circumstances would have to be very peculiar. On the significance of this sort of assumption, see the end of section 6.

Can one dissolve this particular problem by assuming that h and a are independent, and so that $P(a|h) = P(a)$? The assumption of independence is used in this way by Dorling (1979) and Howson and Urbach (1993) to help simplify the standard Bayesian solution to the Quine-Duhem problem, but in fact it is entirely unwarranted. If, as is generally assumed in posing the Quine-Duhem problem, h and a have been part of the same research program for some time, they will certainly not be independent. Suppose, for example, that h has been confirmed by many observations using a certain experimental apparatus and that a states that the apparatus functions as it is supposed to. The independence of h and a entails that $P(h|a) = P(h) = P(h|\neg a)$. Thus to assume that h and a are independent is to assume that one's probability for h would be unaffected by the discovery that the apparatus that has confirmed h so many times has been malfunctioning.

multiplier are illustrated in figure 1. In general, when a conjunct is falsified,




Figure 1: The ratio of the partial posterior $Q(h)$ to the prior $P(h)$, as a function of $P(a|h)$, for three values of $P(h)$. From top to bottom, $P(h) = 0.7, 0.5, 0.3$. In each case, as $P(a|h)$ increases, the multiplier for $P(h)$ decreases.

suspicion will be concentrated on those conjuncts that were more dubious to begin with.

4. Confirming and Disconfirming Evidence

A complete solution to the Quine-Duhem problem must provide a method for assessing the impact of evidence e on a hypothesis h when e neither proves nor refutes ha , but merely confirms or disconfirms it. In this section I provide an appropriate Bayesian method.

As before, it is not sufficient to calculate $P^+(h)$, that is, $P(h|e)$. What is wanted is a procedure that takes into account *only* the fact that e confirms

ha , that is, only the fact that e raises the probability of ha from $P(ha)$ to $P^+(ha)$. (To save words, I assume throughout this section that e raises the probability of ha , but the method I offer applies also to cases where e lowers the probability.)

The standard Bayesian method for taking into account a probability shift that is not a shift to zero or one is Jeffrey conditionalization (Jeffrey 1983).⁸ In general, when the probability of some proposition e changes from $P(e)$ to $P^+(e)$, Jeffrey conditionalization assesses the impact of the change on the probability of a hypothesis h as follows:

$$P^+(h) = P(h|e)P^+(e) + P(h|\neg e)P^+(\neg e)$$

This formula, the natural extension of Bayesian conditionalization, is perfectly suited to the problem at hand.

When the probability of ha goes from $P(ha)$ to $P^+(ha)$ (in virtue of conditionalization on e), the impact on h , represented by the partial posterior $Q(h)$, may be obtained by way of Jeffrey conditionalization as follows:

$$\begin{aligned} Q(h) &= P(h|ha)P^+(ha) + P(h|\neg(ha))P^+(\neg(ha)) \\ &= P^+(ha) + P(h|\neg(ha))P^+(\neg(ha)) \end{aligned}$$

8. Normally, Jeffrey conditionalization is used only in cases where a probability changes “exogenously”, that is, due to some process other than Bayesian conditionalization. The way it is used here is, like the notion of partial posterior probability, slightly unorthodox, but there is no other way for a Bayesian to answer the proponent of the Quine-Duhem thesis.

(It will be noted that the formula used in section 3 is the special case of this formula that results from setting $P^+(ha)$ to zero.)

Now I provide a “posterior objectivist” justification of this more general formula for determining $Q(h)$, that is, I show that the value of δ_{qd} determined by the formula has a number of objectively desirable properties.

I note two important properties. First, when conditionalization on e increases the probability of ha , $Q(h)$ is greater than $P(h)$, hence δ_{qd} is positive. Similarly, when conditionalization on e decreases $P^+(ha)$, δ_{qd} must be negative. The argument is as follows. The prior probability of h , $P(h)$, can be rewritten:

$$\begin{aligned} P(h) &= P(h|ha)P(ha) + P(h|\neg(ha))P(\neg(ha)) \\ &= P(ha) + P(h|\neg(ha))P(\neg(ha)) \end{aligned}$$

Let $k = P(h|\neg(ha))$. (Thus k is equal to what the partial posterior for h would be were ha to be falsified; see section 3.) So the probability of h before e increases the probability of ha is

$$P(h) = P(ha) + kP(\neg(ha))$$

and the partial posterior probability of h after e has increased the probability of ha is

$$Q(h) = P^+(ha) + kP^+(\neg(ha)).$$

Suppose that, on receipt of e , the probability of ha goes up by an amount δ . Then the probability of $\neg(ha)$ goes down by δ . Thus, as a result of the

confirmation of ha , the probability of h goes up by δ and goes down by $k\delta$.

In symbols:

$$\begin{aligned}\delta_{qd} &= Q(h) - P(h) \\ &= (P(h) + \delta - k\delta) - P(h) \\ &= (1 - k)\delta\end{aligned}$$

Since δ is positive and k is less than one (in all but the extreme cases), δ_{qd} is positive. The confirmation of ha increases the probability of h . When ha is disconfirmed, δ is negative, and δ_{qd} is therefore negative.

Second, when the probability of ha changes as a result of e , the magnitude of the resulting impact on h (that is, the absolute value of δ_{qd}) varies negatively with $P(h)$ and positively with $P(a)$. For example:

- If the probability of ha increases, the resulting increase in the probability of h will be greater if $P(h)$ is low than if it is high.
- If the probability of ha increases, the resulting increase in the probability of h will be greater if $P(a)$ is high than if it is low.
- If the probability of ha decreases, the resulting decrease in the probability of h will be greater if $P(a)$ is high than if it is low (a result already seen in section 3 in the extreme case where the probability of ha goes to zero).

This result follows from the same line of thinking as the last result. When e increases the probability of ha , the resulting increase in the probability of h is $(1 - k)\delta$. The increase is greater when k is smaller. (If the probability of ha decreases, so that δ is negative, then the probability of h decreases more when k is smaller.) As shown in section 3, and again in figure 2, k tends




Figure 2: The value of $1 - k$, as a function of $P(a|h)$, for three values of $P(h)$. From top to bottom, $P(h) = 0.3, 0.5, 0.7$. As $P(a|h)$ increases, $1 - k$ increases, and so the magnitude of the impact of the evidence (positive or negative) on the probability of h increases.

to be smaller when $P(h)$ is small or $P(a)$ is large (but see note 7). So the magnitude of δ_{qd} varies negatively with $P(h)$ and positively with $P(a)$.

Some aspects of this result are not surprising. It is a desirable property of any theory of confirmation, for example, that the more highly confirmed a theory, the less affected it is by one more piece of confirming evidence. Much more interesting is the way in which the impact of e on h depends on $P(a)$. What I have shown is that theories that depend on high probability

auxiliaries are relatively much more affected by evidence, confirmatory or disconfirmatory, than those that depend on low probability auxiliaries. A gratifying consequence of this is that, if one wants to learn as much as possible about h from the evidence (in the sense of having the evidence impact one's probability for h by as much as possible), one ought to seek out highly confirmed auxiliary hypotheses. One ought to use equipment that is known to be reliable in conditions where one can be sure that one's background assumptions obtain. In other words, one should proceed in accordance with all the platitudes of experimental science.

5. Beyond the Quine-Duhem Problem

The Quine-Duhem problem is, in many ways, the central problem concerning the role of auxiliary hypotheses in science. One might hope, then, that a Bayesian solution to the Quine-Duhem problem would provide answers to many other questions involving auxiliary hypotheses. My solution cannot be directly employed in a Bayesian treatment of other problems in confirmation theory, however, because it provides a formula for what I call a partial posterior probability rather than for the posterior probability itself. Nevertheless, in those cases where the posterior $P^+(h)$ is close to the partial posterior $Q(h)$ —those cases where the change in the probability of h is almost entirely due to the falsification of ha —the partial posterior can be used as an estimate of the posterior. The results derived in the previous two sections can then be

applied to the new problem. If the problem is solved in a satisfying way, the solution will then provide an argument, in the posterior objectivist style, for Bayesianism itself. In this section I sketch the circumstances in which the partial posteriors provide a good estimate of the posteriors.

At first it may seem quite unlikely that the partial posterior will provide a good estimate of the posterior probability, for the following reason. Suppose that a_1, \dots, a_n are the (mutually exclusive, exhaustive) alternatives to the auxiliary hypothesis a . Then

$$P^+(h) = P^+(ha) + P^+(ha_1) + \dots + P^+(ha_n)$$

The partial posterior represents the change in $P^+(h)$ due to e 's impact on the probability of ha , but ignores the change due to e 's impact on ha_1 , ha_2 , and so on. Unless the net effect of e on the probabilities of these alternative possibilities is close to zero, the partial posterior will therefore differ substantially from the true posterior. It is unlikely that e will have negligible impact on all of ha_1, \dots, ha_n , since at least some of these theories will presumably make predictions about e . Approximate equality of the partial and true posteriors, then, would appear to require that the impact of e on ha 's rival theories approximately balance out. The reader might reasonably think that this is a condition not likely to be satisfied.

This turns out to be wrong, at least in an important group of cases, as I will show below. But rather than immediately pursuing this line of argument, which concerns only a sufficient condition for the partial and true posteriors

to differ, let me state an illuminating necessary and sufficient condition for the two to differ. This will allow a more general treatment of the problem. The condition is as follows. The partial posterior and the true posterior for h will differ just in case

$$P(e|h\neg a) \neq P(e|\neg(ha))$$

Why is this true? From the previous section, if e impacts on h only by way of ha (so that $P^+(h) = Q(h)$), then

$$\begin{aligned} P^+(h) &= Q(h) \\ &= P^+(ha) + P(h|\neg(ha))P^+(\neg(ha)) \end{aligned}$$

The converse also holds. By the mathematics of probability alone,

$$P^+(h) = P^+(ha) + P^+(h|\neg(ha))P^+(\neg(ha))$$

Thus

$$P^+(h) = Q(h) \quad \text{iff} \quad P(h|\neg(ha)) = P^+(h|\neg(ha))$$

$P^+(h)$ will differ from $Q(h)$, then, whenever e causes a change in the probability of $P(h|\neg(ha))$. Since

$$P(h|\neg(ha)) = \frac{P(h\neg a)}{P(\neg(ha))}$$

the partial and true posteriors will differ just in case

$$\frac{P(h\neg a)}{P(\neg(ha))} \neq \frac{P^+(h\neg a)}{P^+(\neg(ha))}$$

This happens just in case e impacts differently on the probability of $h\neg a$ than on the probability of $\neg(ha)$, that is, by Bayes' theorem, just in case $P(e|h\neg a)$ is not equal to $P(e|\neg(ha))$.

There are three conditions under which $P(e|h\neg a)$ will be different from $P(e|\neg(ha))$ (note that only in the third do the alternatives to a , disjoined to make $\neg a$, play an essential role):

1. The central hypothesis h , without the help of any other hypothesis, assigns a probability to e not equal to $P(e|\neg(ha))$.
2. The central hypothesis h , in conjunction with some set of mutually exclusive, exhaustive hypotheses v_1, v_2, \dots that are not connected with a , assigns a probability to e on average⁹ not equal to $P(e|\neg(ha))$. This is the case in the Newstein scenario (see section 2.4), where the hypotheses v_1 and v_2 are “Newstein is reliable” and “Newstein is not reliable”. (Newstein’s predictions, being known for certain, are here regarded as background information, thus implicitly taken into account; see note 5.)
3. The central hypothesis h , in conjunction with the various alternatives to a , assigns a probability to e on average not equal to $P(e|\neg(ha))$. This is most likely to occur in real life when both (a) there are a limited range of alternatives to a and (b) each of the alternatives, in conjunction

9. Here and below it should be understood that I mean an average weighted by the prior probabilities of the hypotheses v_i in conjunction with h , as they appear in the relevant version of the theorem of total probability.

with h , assigns a well-defined probability to e . Then there will be a well-defined probability $P(e|h\neg a)$ that may differ from $P(e|\neg(ha))$. An example of such a situation is a case where a concerns only the value of some physical constant that appears (without being assigned a specific value) in h .

Suppose that a is a group of auxiliary hypotheses of the sort needed to get predictions from a central hypothesis h . Will any of conditions one to three apply, or will $P^+(h)$ and $Q(h)$ be approximately equal? Condition one, according to Duhem's thesis, will seldom apply. The question whether condition two applies will have to be determined on a case by case basis, but I assume that, in many cases, the effect will not be significant. (Note that, unlike proponents of the standard Bayesian solution to the Quine-Duhem problem, I am able to give a formal condition for there being no Newstein effect. This is possible only because I have at my disposal a formal definition of δ_{qd} .)

It is after all, then, the alternatives to a that pose the most immediate threat to the partial posterior's reliability as an estimate of the true posterior. As earlier promised, I will present an argument that condition three is satisfied in a certain important class of cases, so that the effect of e on the different rival theories ha_1, \dots, ha_n cancels out in the required sense.

The class of cases I have in mind is that in which a concerns the experimental conditions under which the evidence is produced, that is, where a is a conjunction of a number of hypotheses concerning such things as the state

of the equipment, the background conditions for the experiment, and so on. Suppose that, in such a case, e is something like “The readout says 5.17”. Consider two ways that a might turn out to be false.

First, the apparatus might be malfunctioning. Then the probabilities for different figures on the readout are determined by the probabilities for the different ways the apparatus might fail. The probability of e —of obtaining the particular reading that was obtained—conditional on the apparatus being broken will, in most cases, be very nearly zero.¹⁰

Second, an assumption about the background conditions might be false. Then the probabilities for different figures on the readout are determined by the probabilities for all the possible background conditions. The probability of e conditional on the background assumptions being false is also likely to be very nearly zero.

What if h is false? The many different alternatives to h presumably probabilify many different values for the readout (including a blank, when the experiment fails completely). Only a very small proportion of these alternatives will, in conjunction with some given set of auxiliary hypotheses, predict e . So the probability of e , given that h is false, is very nearly zero.

Bringing the possibilities together, the probability of e conditional on

10. I say “in most cases” because of the following sort of exception. If it is known that a broken apparatus typically produces a certain value on the readout (say zero), and if that value appears, then the argument just given does not go through: the probability of e conditional on equipment failure is not, after all, close to zero. There will be other such exceptions for the same kinds of reasons.

(a) the apparatus being broken, (b) the background assumptions being false, (c) h being false, or any combination of these, will be very nearly zero. If all these probabilities are nearly equal, then the probability of e given $h \rightarrow a$ is going to be nearly equal to the probability of e given $\neg ha$ or $\neg h \rightarrow a$, and so condition three will be satisfied. That the three probabilities are close to zero is not, however, sufficient to conclude that they are nearly equal in the relevant sense. The reason is that what matters is the ratio of the probabilities. That two probabilities are close to zero does not mean that one is not many times larger than the other. To reach my conclusion, then, I must show that the ratio of the nearly zero probabilities is near one.

At this point, I appeal to the fact that we are dealing with the subjective probabilities of flesh and blood scientists. For the value of, say, $P(e|h \rightarrow a)$ to be twice that of $P(e|\neg ha)$, when both probabilities are very nearly zero, requires that both probabilities be well defined to a number of decimal places. (How many decimal places depends, more or less, on the accuracy of the measuring apparatus.) But in the sort of case under consideration, it is highly unlikely that scientists will have such precise subjective probabilities. Ask a working scientist what is the chance of the readout showing exactly 5.17, given that the apparatus is broken, and they will reply “very nearly zero”. Ask them what is the chance of the readout showing exactly 5.17, given that the central hypothesis is false, and they will reply “very nearly zero”. Further, and most important of all, were one to ask the scientist what the relative size of the two nearly zero probabilities is, they would likely

reply, I suggest, “about the same”. This, I submit, reflects just about the full extent of the scientist’s opinions on the matter.

Ought the scientist to try harder? Not at all. The cost, in time and energy, of enumerating the various alternatives to h and a , and then of fishing up from the well of consciousness subjective probabilities for e given each of these different alternatives, clearly far outweighs the additional precision one might hope to gain in one’s subjective probabilities for h and a . Trying harder, then, would be irrational: the scientist should stick with their initial estimate that the near zero probabilities are “about the same”. All of this is enough to satisfy condition three.¹¹

I conclude that it is both psychologically accurate to suppose that scientists will not, and in accordance with the canons of procedural rationality to hold that they ought not, go to the trouble of assigning extremely precise values to $P(e|h\rightarrow a)$ and $P(e|\neg(ha))$. Thus, in the kind of circumstances sketched above, we may assume, and think it reasonable, that scientists set a value for $P(e|h\rightarrow a)$ that is pretty much equal to $P(e|\neg(ha))$.

It follows that there is good reason for thinking that, in many cases where a is a realistic brew of the kinds of assumptions needed to get predictions

11. There are, of course, several different ways of formalizing the imprecision in subjective probabilities of this sort. One may give probability intervals rather than exact probabilities. In the case above, I suggest, the intervals for the two conditional probabilities will be about the same. Or one may allocate probability in discrete units, for example, multiples of 0.001. In the case above, I suggest, the two conditional probabilities will take values equal to a single multiple of the base unit, the numerical equivalent of “very nearly zero”.

from a central hypothesis, the partial posterior $Q(h)$ is a good estimate of the posterior $P^+(h)$. In these cases, then, posterior probabilities obtained by Bayesian conditionalization will behave in a certain way, namely, in the way the partial posteriors were shown to behave in sections 3 and 4. This sort of result allows some very strong conclusions to be drawn about the properties of Bayesian conditionalization, conclusions which are put to use in the next section.

6. Ad Hoc Auxiliary Hypotheses

When a theory ha is falsified by some piece of evidence e , the central hypothesis h is rescued from falsification when there exists (as there almost always does) an alternative to a , say a' , such that ha' is not inconsistent with e . When a' is picked out specifically for its ability to rescue h from e , it is called ad hoc. Often the forced move to an ad hoc auxiliary hypothesis is thought to seriously damage the central theory's prospects. But this is not always true: the classic example of a "rescue" that did not seriously affect the central theory's credibility is the postulation of the existence of the planet Neptune to accommodate perturbations in the orbit of Uranus not otherwise consistent with Newtonian theory. Not only did this rescue not seriously affect the credibility of Newtonian theory, it vastly increased the credibility of the rescuing ad hoc hypothesis, the hypothesis that an eighth planet existed (and had a certain orbit). Howson and Urbach (1993, 156)

quote John Herschel's confident declaration that the "rescue" in fact constituted a discovery: Uranus's existence was known "with a certainty hardly inferior to that of ocular demonstration".

This raises a problem: when does the rescue of a central hypothesis h by an ad hoc hypothesis a' seriously dent the credibility of h (a *desperate rescue*), and when does it greatly increase the probability of a' (a *glorious rescue*)? Howson and Urbach (1993, 157–8) give a Bayesian answer to this question. A rescue by a' is glorious when the probability of a' conditional on e is high; it is desperate when that probability is low. The reason is that, when $P(a'|e)$ is high, $P^+(a')$ is high, and so $P^+(ha')$ will tend to be reasonably high, while when $P(a'|e)$ is low, $P^+(a')$ is low and so $P^+(ha')$ will have to be at least as low. A high $P^+(ha')$ redounds to the credit of h and a' ; a low $P^+(ha')$ embarrasses them. (See also Horwich 1982, 105–8.)

This answer to the question is, I think, at best incomplete. Howson and Urbach's story perhaps explains the glory or disgrace that accrues, after a rescue, to the research program as a whole (which I assume is proportional to $P^+(ha')$), but it does not answer two important questions:

1. Why, in a glorious rescue, does the probability of a' increase? On Howson and Urbach's story, we learn nothing about the dynamics of the probability of a' at all. Insofar as explaining the glory of a glorious rescue involves explaining why the probability of a' , conditional on the evidence, is so high, Howson and Urbach might be thought to be assuming part of that which was to be explained.

2. Why, in a desperate rescue, does the probability of h decrease? On Howson and Urbach's story about a desperate rescue, the low posterior probability of ha' would appear to be explained entirely by the low posterior probability of a' ; the low probability of ha' , then, cannot account for the decrease in the probability of h .

In what follows, I hope to do better. I will assume that the relevant posteriors are approximately equal to the partial posteriors, for the reason that conditions one, two and three of section 5 hold. The discovery of Neptune provides a good example of a case where condition three, in particular, is satisfied, that is, where the central hypothesis h , in conjunction with the various alternatives to a , assigns a probability to e on average equal (because very nearly zero) to $P(e|\neg(ha))$. To see this, consider the following. The probability of Uranus's orbit's being perturbed in exactly the observed way, given that a is false, that is, given that there are more planets, will be very nearly zero, since only a vanishingly small proportion of the possible configurations of any extra planets would give rise to the observed perturbations. Similarly, the probability of Uranus's orbit's being perturbed in exactly the observed way, given that h is false, that is, given that some theory of celestial mechanics other than Newton's is true, will also be very nearly zero, since only a vanishingly small proportion of the possible rivals to Newton's theory (not that there was any known, plausible rival that explained the evidence) would give rise to the observed perturbations. It is a safe bet, I think, that virtually no scientist at the time would have held that one of these probabilities was

higher than the other. Further, to do so would, I think, have been rightly regarded as rather unreasonable. Thus section 5's condition three may, along with the other conditions, be assumed to have held in this case.

Given the approximate equality of the partial and true posterior probabilities, the solution to the Quine-Duhem problem presented in this paper provides a superior alternative to (or if you prefer, a completion of) Howson and Urbach's solution. Recall that, when ha is falsified, h suffers more when $P(h)$ is smaller and $P(a)$ is larger (see figure 1). If $P(h)$ is very high, as the probability for Newtonian theory was in 1846 (the year of Herschel's pronouncement), it will decrease very little upon the falsification of ha . Central hypotheses with lower probabilities will suffer more from the refutation of ha ; their rescue will be accompanied by a more embarrassing probability loss.

The impact of the falsification of ha on a is determined in the same way: a suffers more when its probability is low or the probability of h is high. When the probability of h is near one and that of a is not, the probability of a will suffer badly. As a result, the probabilities of the rivals to a will increase, provided that those rivals, in conjunction with h , probabilify e . Since the ad hoc hypothesis a' is constructed especially to combine with h so as to probabilify e , its probability will rise. If it is by far the least unlikely such hypothesis (as was the case in the discovery of Neptune), it will gain nearly all of the probability lost by a . Thus its credibility will increase dramatically.

In summary, a glorious rescue occurs roughly when $P(h)$ is considerably higher than $P(a)$, while a desperate rescue occurs roughly when $P(h)$ is

considerably lower than $P(a)$. In words, a glorious rescue occurs when the auxiliary hypothesis receives most of the blame for a false prediction, and is rightly discarded by researchers in favor of some other auxiliary hypothesis that makes the correct prediction. (The degree of glory, I remark in passing, is perhaps inversely proportional to the prior probability of the ad hoc hypothesis.) A desperate rescue occurs when the central hypothesis receives most of the blame for a false prediction, but where researchers cling to the central hypothesis and discard the evidently superior auxiliary.

The reader will note that the calculation of the effect of the falsification of ha on the probabilities of h and a does not take into account the details of the falsifying evidence. That is, the question of how ha was falsified makes no difference to the impact of the falsification on h and a . Is this in general true? It is not. According to Bayesian confirmation theory, the details of the falsification can have any kind of impact on h and a that you like (by way of Newstein effects and so on). In the last section, however, I argued that in certain, very common circumstances, the posteriors and the partial posteriors are roughly equal. This is equivalent to showing that, in those same circumstances, the details of the falsification do not matter. Because the circumstances obtain in the Neptune case, then, the details of the perturbations in the orbit of Uranus have a negligible effect on the fate of h and a . (Of course, these details most definitely are taken into account in explaining the fact that a' , and not some other alternative to a , was the chief beneficiary of a 's loss of probability.) More generally, if the conditions stated

in the last section hold, then the details of the manner in which a theory is falsified make no difference, on the Bayesian account, to whether an ad hoc rescue of the theory's central hypothesis is desperate or glorious.

Let me conclude this section and this paper by suggesting that my proposed Bayesian treatment of the Uranus problem, and more generally, my proposed Bayesian treatment of the problem of grounding the distinction between desperate and glorious rescues, endows the confirmation process with many of the properties that one might reasonably desire in the cases to which the treatment applies. As such, the treatment provides the basis for a part of the posterior objectivist case for Bayesian confirmation theory: Bayesianism is a good confirmation theory because it reproduces, and indeed explains, many of the properties that confirmation is known to have.

It should be noted, however, that Bayesianism has been shown to have desirable properties only given certain assumptions about the priors (in particular, those assumptions made in note 7 and section 5). These assumptions are, I think, reasonable and realistic, but I have offered no proof that they are a priori justified. To understand the circumstances in which the assumptions would be justified is, I think, an important and distinct project for Bayesianism. This is, in my opinion, the way forward for Bayesian confirmation theory: not to show that certain priors are a priori justified, nor to prove that Bayesian conditionalization will produce desirable results regardless of the priors, but to show that given certain reasonable and realistic assumptions about the priors, Bayesian conditionalization will produce de-

sirable results. This sort of Bayesianism will, it must be admitted, lack the philosophical grandeur of a priori epistemology, but the reward for modesty is feasibility.

References

- Dorling, J. (1979). Bayesian personalism, the methodology of scientific research programmes, and Duhem's problem. *Studies in History and Philosophy of Science* 10:177–87.
- Duhem, P. (1954). *The Aim and Structure of Physical Theory*. Translated by P. P. Wiener. Princeton University Press, Princeton, NJ.
- Earman, J. (1992). *Bayes or Bust?* MIT Press, Cambridge, MA.
- Edwards, W., H. Lindman, and L. J. Savage. (1963). Bayesian statistical inference for psychological research. *Psychological Review* 70:193–242.
- Horwich, P. (1982). *Probability and Evidence*. Cambridge University Press, Cambridge.
- Howson, C. and P. Urbach. (1993). *Scientific Reasoning: The Bayesian Approach*. Second edition. Open Court, Chicago.
- Jaynes, E. T. (1983). *Papers on Probability, Statistics, and Statistical Physics*. D. Reidel, Dordrecht. Edited by Roger Rosenkrantz.
- Jeffrey, R. C. (1983). *The Logic of Decision*. Second edition. Chicago University Press, Chicago.
- Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, volume 2. University of California Press, Berkeley, CA.
- Rosenkrantz, R. D. (1981). *Foundations and Applications of Inductive Probability*. Ridgeview, Atascadero, CA.
- Strevens, M. (1999). Objective probabilities as a guide to the world. *Philosophical Studies* 95:243–75.

Worrall, J. (1994). Falsification, rationality, and the Duhem problem: Grünbaum versus Bayes. In J. Earman, A. I. Janis, and G. J. Massey (eds.), *Philosophical Problems of the Internal and External Worlds: Essays on the Philosophy of Adolf Grünbaum*. University of Pittsburgh Press, Pittsburgh.