

Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions

JAMES HAWTHORNE (forthcoming in *Mind*)

I argue that Bayesians need two distinct notions of probability. We need the usual degree-of-belief notion that is central to the Bayesian account of rational decision. But Bayesians also need a separate notion of probability that represents the degree to which evidence supports hypotheses. Although degree-of-belief is well suited to the theory of rational decision, Bayesians have tried to apply it to the realm of hypothesis confirmation as well. This double duty leads to the problem of old evidence, a problem that, we will see, is much more extensive than usually recognized. I will argue that degree-of-support is distinct from degree-of-belief, that it is not just a kind of counterfactual degree-of-belief, and that it supplements degree-of-belief in a way that resolves the problems of old evidence and provides a richer account of the logic of scientific inference and belief.

1. Introduction

I contend that Bayesians need two distinct notions of probability. We need the usual *degree-of-belief* notion that is central to the Bayesian account of rational decision. But Bayesians also need a separate notion of probability that represents the degree to which evidence *supports* hypotheses. Although *degree-of-belief* is well suited to the theory of rational decision, Bayesians have tried to apply it in the realm of hypothesis confirmation as well. This double duty leads to the problem of old evidence, a problem that, we will see, is much more extensive than usually recognized.

Degree-of-support is closely akin to the notion of logical probability. That notion has gotten short shrift because of the failure of attempts like those of Keynes and Carnap to explicate it in terms of syntactic structure. But unlike logical probability, *degree-of-support* is not a purely formal notion. Its logical nature is most evident in the Bayesian treatment of the *likelihoods*, the conditional probabilities that express what hypotheses or theories *say* about evidential claims. I will argue that *degree-of-support* is distinct from *degree-of-belief*, that it is not just a kind of counterfactual *degree-of-belief*, and that it supplements *degree-of-belief* in a way that resolves the problems of old evidence and provides a richer account of the logic by which scientific hypotheses should be evaluated.

My argument basically runs as follows:

1. Bayesian accounts of how evidence supports hypotheses in the sciences usually rely on the objectivity or ‘publicness’ of the likelihoods that occur in Bayes’s theorem.
2. However, old evidence problems show that the numerical values taken by the likelihoods in an agent’s *degree-of-belief* function *must* often deviate considerably from the objective or public values that likelihoods are supposed to have.
3. Thus, if a Bayesian account is to model both the agent’s *belief strengths* and the role of objective likelihoods in the evidential *support* of hypotheses in of scientific contexts, and if such *support* is supposed to influence the agent’s *belief strengths*, then the account will have to draw on two separate probability functions, and it will need to tell us how the agent’s *support function* is supposed to influence his *belief function*.

Finally, I will describe how, on a two-function model of scientific inference, Bayesian *support* should inform Bayesian *belief*.¹

¹ The idea that Bayesian epistemology should draw on two distinct probability functions is not new. Carnap (1971) suggested a two-function Bayesian model. He calls the *degree-of-belief* notion ‘rational credence’ and calls the *degree-of-support* notion ‘credibility’. He takes *initial credence functions* to derive from *credibility functions*, which should themselves be *logical probability functions*. Brian Skyrms largely adopts this Carnapian idea in the third edition of *Choice and Chance* (1986, Ch. 1, Ch. 6, Sects. 7 and 8), but does not identify his version of *credibility functions* with Carnapian *logical*

Let me elaborate a bit. The Bayesian account of how evidence supports hypotheses in scientific contexts often relies on the objectivity or ‘publicness’ of the likelihoods that occur in Bayes’s theorem. The likelihoods are supposed to represent the empirical content of hypotheses, what hypotheses *say* about the evidence. The principal reason that Bayesians appeal to Bayes’s theorem to evaluate the support hypotheses receive from evidence is that Bayes’s theorem expresses evidential support in terms of objective, or ‘public’, or ‘better known’ likelihoods. In section 2 I describe the Bayesian account of evidential support in a way that emphasizes the centrality of the likelihoods. In section 3 I argue that the objectivity or publicness of likelihoods is vital to the kind of scientific objectivity articulated by the Bayesian account of confirmation.

Then, in section 4 I show how old evidence problems undermine the objectivity of the likelihoods as they are expressed by agents’ *degrees-of-belief*. I show that belief function likelihoods should almost inevitably deviate from the objective or public values that likelihoods are often supposed to have in Bayesian confirmational contexts. For, belief function likelihoods represent how strongly the agent is supposed to believe an evidence claim when a given hypothesis is added to *everything else the agent knows*. So, when the agent is already certain of an evidence statement, her *belief-function* likelihoods for that statement must be 1 on every hypothesis. Moreover, I describe new versions of the old evidence problem that arise in cases in which the agent is uncertain of the truth of an evidence claim but possesses some bit of information relevant to its truth. Agents will almost always possess such information, usually of a quite trivial kind. I show that when they do, their belief function likelihoods should almost surely deviate from objective or public values. Thus, if objective or public likelihoods are to maintain their role in Bayesian accounts of evidential support, these likelihoods must be part of a Bayesian probability function that is distinct from the agent’s *belief function*. I call this second probability function a *support function*. I also argue that support functions are not some sort of *counterfactual belief functions*. Rather, they are more closely akin to a certain conception of logical probability.

Bayesian *belief functions* represent agents’ belief-strengths, which are the probabilities they use in decision-making. However, in scientific contexts *support functions* play an important role in supporting or justifying the agent’s belief strengths for hypotheses. In section 5 I argue this point further. Thus, a more complete Bayesian model of scientific inference should employ a distinct probability function for each notion and describe how evidential support should influence belief. In section 6 I flesh out how this should work.

2. Bayesian induction

In this section I will give an overview of Bayesian confirmation that emphasizes the central role of the likelihoods. This will mostly be old hat to aficionados; but I want to make sure we are all on the same page. And the reader may find a few unfamiliar wrinkles. So please bear with me.

Let α be a Bayesian agent and let P_α be a probability function that represents for her the degree to which some statements support others. I prefer to think of P_α as a semantic conditional probability function defined on a language for predicate logic — in particular, I prefer to treat them formally as Popper-functions, as described by Hartry Field (1977). But what I’ll say here doesn’t much depend on this. So let P_α be defined on sentences or propositions in your favorite Bayesian way. And although I’ll refer to P_α as a *degree-of-support function*, you may for now think of it as one of the usual Bayesian *degree-of-belief functions*, if you wish.

probabilities. (Skyrms’s fourth edition moves the chapter 1 material to chapter 2 and completely drops the relevant chapter 6 material.) Skyrms calls the *degree-of-belief* notion ‘epistemic probability’, and calls the *degree-of-support* notion ‘inductive probability’. More recently Marc Lange (1999) also argues for a two-function Bayesian model. On Lange’s view support functions provide a basis for justificatory arguments for the agent’s belief strengths. The two-function model I’ll present resembles Lange’s and Skyrms’s approaches in certain respects, but also differs in important ways. I’ll say more about this later.

Let $\{H_1, H_2, \dots\}$ be a mutually exclusive and exhaustive set of hypotheses or theories about some subject matter. It may be finite or infinite. If finite, it may contain a *catch-all hypothesis*, H_K that says that none of the other hypotheses are true: $(\neg H_1 \dots \neg H_m)$. Let ‘ C ’ represent a description of observational or experimental Conditions and let ‘ E ’ represent a description of an *Evidential* outcome or observation that may result. In addition, let ‘ B ’ represent explicit *Background* information and auxiliary hypotheses that are not at issue in the assessment of hypotheses H_i , but mediate the connection between hypotheses and evidence. When H_i is deductively related to the evidence, B supplies whatever auxiliaries mediate the logical entailments: $H_i \cdot B \cdot C \models E$ or $H_i \cdot B \cdot C \models \neg E$. When H_i is related to the evidence statistically rather than deductively (either because H_i is explicitly statistical or because an auxiliary statistical hypothesis connects H_i to the evidence), B supplies whatever auxiliaries mediate the so-called *direct inference likelihoods* of the evidence on the hypothesis: $P[E | H_i \cdot B \cdot C] = r$. Although presentations of Bayesian inference often suppress C and B , they are an important part of the logic of the likelihoods. So I’ll continue to make them explicit.²

A *likelihood* is a probability that expresses how likely the evidence is on a given hypothesis. If a hypothesis together with auxiliaries and observation conditions deductively entails the evidence, the axioms of probability make the likelihood completely objective: $P[E | H_i \cdot B \cdot C] = 1$ if $H_i \cdot B \cdot C \models E$; $P[E | H_i \cdot B \cdot C] = 0$ if $H_i \cdot B \cdot C \models \neg E$. For statistical direct inferences, where $H_i \cdot B$ specifies a statistical hypothesis that bears directly on E via C , most Bayesians also take the likelihoods ‘ $P[E | H_i \cdot B \cdot C] = r$ ’ to be highly objective. Indeed, some Bayesian logicians have attempted to spell out the logic of statistical direct inferences in terms of the logical forms of the sentences involved.³ But regardless of whether statistical direct inferences can be formalized, Bayesians tend to consider them and many of the other likelihoods that arise in scientific contexts to have highly objective or inter-subjectively agreed values. That is, although various Bayesian agents, α, β , etc., may widely disagree about the plausibility of various hypotheses, they tend to largely agree on the likelihoods. (I’ll say much more about this in the next section.) I will mark the high degree of objectivity or inter-subjective agreement on the values of the likelihoods by dropping the subscript ‘ α ’ from expressions that represent them.

One more wrinkle before looking closely at Bayes’s Theorem. Evidence typically accumulates over time. That is, hypotheses are tested by a sequence of experiments or observations, C_1, C_2, \dots, C_n , that result in corresponding outcomes E_1, E_2, \dots, E_n . I will abbreviate the conjunction of the first n experimental or observation conditions as ‘ C^n ’ and the conjunction of their outcomes as ‘ E^n ’.

Now let us take a look at several forms of Bayes’s Theorem that Bayesians rely on to capture the logic of confirmation. The simplest form is this:

$$(1) P_\alpha[H_i | B \cdot C^n \cdot E^n] = P[E^n | H_i \cdot B \cdot C^n] \cdot P_\alpha[H_i | B \cdot C^n] / P_\alpha[E^n | B \cdot C^n].$$

Equation (1) expresses the *posterior probability* of H_i on evidence and background in terms of the *likelihood* of the evidence on the hypothesis, the *prior probability* of the hypothesis on background and observation conditions, and the *simple probability* of the evidence on background and observation conditions. This latter probability is sometimes called the *expectedness* of the evidence. Typically, the *prior* probability of the hypothesis and the *expectedness* of the evidence are highly subjective. Bayesians are often willing to accept the subjectivity of the prior probabilities. But they find the subjectivity of the *expectedness* more troubling, and seek a remedy.

One way to avoid the subjective *expectedness* of the evidence is to consider a ratio form of Bayes’s theorem, a form that compares hypotheses one pair at a time:

² Background information may, at least in part, consist of auxiliary hypotheses that are themselves subject to confirmation on some additional body of evidence. If we can treat each H_i as a super-hypothesis that contains all auxiliaries it needs to speak to the evidence, then just let ‘ B ’ be a tautology.

³ These attempts have not been completely satisfactory. But I think the project can succeed. (I hope to offer a specific proposal elsewhere.) The series of papers by Levi (1977), Kyburg (1978), and Levi’s response (1978) provides an illuminating discussion of the logic of direct inference and the difficulties involved in providing a formal treatment. See Levi (1980) and Harper (1981) for sophisticated proposals.

$$\begin{aligned}
 (2) \quad \frac{P_a[H_j | B \cdot C^n \cdot E^n]}{P_a[H_i | B \cdot C^n \cdot E^n]} &= \frac{P[E^n | H_j \cdot B \cdot C^n]}{P[E^n | H_i \cdot B \cdot C^n]} \cdot \frac{P_a[H_j | B]}{P_a[H_i | B]} \cdot \frac{P_a[C^n | H_j \cdot B]}{P_a[C^n | H_i \cdot B]} \\
 &= \frac{P[E^n | H_j \cdot B \cdot C^n]}{P[E^n | H_i \cdot B \cdot C^n]} \cdot \frac{P_a[H_j | B]}{P_a[H_i | B]}
 \end{aligned}$$

assuming the conditions C^n are no more likely on H_i than on H_j — i.e. that $P_a[C^n | H_j \cdot B] = P_a[C^n | H_i \cdot B]$. (This amounts to supposing that the occurrence of the observation conditions does not itself constitute evidence for or against the hypotheses. This supposition may be substantially relaxed without affecting anything I'll say here; but doing so would add inessential complexity.)

Thus, the only really subjective element that contributes to the ratio of posterior probabilities is the ratio of the prior probabilities.

In Bayesian confirmation theory the prior probability of a hypothesis represents how plausible the agent takes the hypothesis to be *prior to* taking the evidence ($C^n \cdot E^n$) into account. Critics point out that such plausibility ratings are usually highly subjective, and claim that these factors completely undermine the objectivity of Bayesian confirmation. In the sciences, however, plausibility assessments are not mere subjective whims. They are often backed by forceful conceptual arguments (e.g., thought experiments). Such arguments tend to be particularly useful in bringing the scientific community into agreement on the *extreme implausibility* of various 'logically possible' alternatives. Agents may often disagree on the relative strengths of plausibility arguments for the remaining, viable hypotheses, and so may disagree widely on prior plausibility assessments. However, such diversity need cause no particular difficulty for the objectivity of hypothesis confirmation, provided that the likelihoods are fairly objective, and that sufficient empirical evidence becomes available. For, Bayesians point out, empirical evidence often tends to override, or 'wash out' the diversity among the plausibility assessments of agents.

To see how 'washing out' may occur, notice that if, as evidence accumulates, the ratios of likelihoods $P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n]$ approach 0, then the posterior probability of H_j must approach 0 as well — the evidence comes to strongly refute H_j . One kind of Bayesian convergence theorem shows that if ($H_i \cdot B \cdot C^n$) is true, these likelihood ratios will *very probably* approach 0 as the evidence accumulates. Thus, by equation (2), the posterior probability of H_j must very probably approach 0 as well, as evidence accumulates. Call this result the *Likelihood Ratio Convergence Theorem*. It is a version of the *Weak Law of Large Numbers*. Let's consider what this theorem says in a bit more detail.

Choose any value of ε , as small as you like. The expression ' $\{E^n : P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n] < \varepsilon\}$ ' represents the set of all statements E^n describing possible sequences of outcomes of the sequence of experiments or observations C^n that would yield a likelihood ratio $P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n]$ with a value less than ε . Now, let the expression ' $\forall \{E^n : P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n] < \varepsilon\}$ ' represent the statement formed by taking the disjunction of all such outcome sequences E^n . The *Likelihood Ratio Convergence Theorem* says that the *likelihood of this disjunctive sentence*,

$$P[\forall \{E^n : P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n] < \varepsilon\} | H_i \cdot B \cdot C^n],$$

must be at least $1 - (q/n)$, where the value of q is an explicitly calculable factor greater than zero.⁴ Thus, the true hypothesis H_i (aided by $B \cdot C^n$) *says* (via a *likelihood*) that it becomes highly likely (as close to 1 as you please) that as evidence increases, one of the possible outcome sequences E^n will occur that yields a likelihood ratio

⁴ In (Hawthorne 2004b, Sect. 5) I show how the value of q may be explicitly computed from ε together with an information theoretic measure of how empirically distinct H_j is from H_i on the possible outcomes of the C_k .

$P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n]$ less than any chosen value of ϵ . As this happens the posterior probability of H_i 's false competitors, H_j , must approach 0.⁵

If we consider the ratio versions of Bayes's Theorem in equation (2) applied to each of H_i 's competitors, and add these together, we get a form of Bayes's Theorem in terms of the *odds* against H_i (where the odds for A given B is defined as $\Omega_\alpha[A | B] = P_\alpha[A | B] / P_\alpha[\neg A | B]$):

$$(3) \quad \Omega_\alpha[\neg H_i | B \cdot C^n \cdot E^n] = \sum_{j \neq i} \frac{P[E^n | H_j \cdot B \cdot C^n]}{P[E^n | H_i \cdot B \cdot C^n]} \cdot \frac{P_\alpha[H_j | B]}{P_\alpha[H_i | B]} + \frac{P_\alpha[E^n | H_K \cdot B \cdot C^n]}{P[E^n | H_i \cdot B \cdot C^n]} \cdot \frac{P_\alpha[H_K | B]}{P_\alpha[H_i | B]} .$$

Notice that if a catch-all hypothesis is needed, the likelihood of evidence relative to it will not generally enjoy the same kind of objectivity as the likelihoods for *specific* hypotheses. I've left the subscript ' α ' on the likelihood for the catch-all to mark its subjectivity.

As new hypotheses are discovered, they are 'peeled off' of the catch-all hypothesis. That is, when a new hypothesis H_{m+1} becomes explicit, the old H_K is replaced by a new catch-all, H_{K^*} , of form $(\neg H_1 \cdot \dots \cdot \neg H_m \cdot \neg H_{m+1})$. The axioms of probability require that the prior probability of this new catch-all hypothesis comes from diminishing the prior of the old catch-all: $P_\alpha[H_{K^*} | B] = P_\alpha[H_K | B] - P_\alpha[H_{m+1} | B]$. So the influence of the catch-all term may diminish towards 0 over time as new alternative hypotheses are made explicit.

If increasing evidence drives the likelihood ratios comparing H_i with each of its competitors towards 0, then the odds against H_i , $\Omega_\alpha[\neg H_i | B \cdot C^n \cdot E^n]$, approaches 0 — provided the priors of catch-all terms, if needed, approach 0 as new hypotheses become explicit and are peeled off. And as $\Omega_\alpha[\neg H_i | B \cdot C^n \cdot E^n]$ goes to 0, the posterior probability of H_i goes to 1. The relationship between the odds against H_i and its posterior probability is this:

$$(4) \quad P_\alpha[H_i | B \cdot C^n \cdot E^n] = 1 / (1 + \Omega_\alpha[\neg H_i | B \cdot C^n \cdot E^n]).$$

The *Likelihood Ratio Convergence Theorem* implies that if $H_i \cdot B \cdot C^n$ is true, then indeed each of the likelihood ratios in equation (3) will very probably approach 0 as the evidence increases.⁶ Thus, if the

⁵ The result described here is one of two parts of the *Likelihood Ratio Convergence Theorem*. It applies to the sub-sequence of the total evidence that satisfies the following condition: for each C_k in sequence C^n , whenever H_j says that outcome O_{kv} is impossible (i.e. $P[O_{kv} | H_j \cdot B \cdot C_k] = 0$), H_i also says that O_{kv} is impossible. The other part of the theorem applies to whatever remaining sub-sequence of the evidence violates this condition. It says: if for some $\epsilon > 0$, a the remaining sub-sequence C^h is made up of individual observations C_k that violate the above condition in that H_j says an outcome O_{kv} is impossible but H_i gives it a likelihood as great as ϵ (i.e., $P[O_{kv} | H_i \cdot B \cdot C_k] \geq \epsilon$), then $P[\bigvee \{E^h : P[E^h | H_j \cdot B \cdot C^h] / P[E^h | H_i \cdot B \cdot C^h] = 0\} | H_i \cdot B \cdot C^h] \geq 1 - (1 - \epsilon)^h$, which approaches 1 as h increases. Thus, for this sub-sequence the true hypothesis H_i says that as evidence of this sort accumulates, it is highly likely (as near 1 as you please) that one of the outcome sequences E^n will occur that yields a likelihood ratio $P[E^n | H_j \cdot B \cdot C^n] / P[E^n | H_i \cdot B \cdot C^n] = 0$, which would result in a posterior probability of 0 for H_j . See Hawthorne (1993) or (2004b, Sect. 5) for more details, including proof of the theorem.

Notice that this convergence theorem does not involve *second-order probabilities* — it does not involve the probability of a probability. Rather, the theorem merely expresses the *likelihood* that some particular disjunctive sentences will be true. This theorem does not require that evidence consists of identically distributed events (like repeated tosses of a coin), it does not draw on countable additivity, and the explicit lower bounds on convergence means that there is no need to wait for the infinite long run for convergence to occur. Indeed, since this result only relies on the likelihoods, it applies even if the Bayesian agent updates his support function from time to time by reassessing the prior plausibilities of hypotheses. (Earman (1992, Ch. 6) provides a thorough analysis of other Bayesian convergence results.)

⁶ This depends, of course, on H_i being empirically distinct from H_j in the sense that for some conditions

conditions for Bayesian convergence are satisfied (and the catch-all, if there is one, diminishes towards 0 through peeling off), then the posterior probability of the true hypothesis will very probably approach 1 as evidence increases.

In any case, regardless of whether some Bayesian convergence theorem applies, equations (2), (3), and (4) show the crucial role played by the likelihoods in the Bayesian assessment of evidential support for scientific hypotheses. They show that if the likelihoods (or, at least, the ratios of likelihoods) are objective or have publicly agreed values, then the only subjective factors in the assessment of the posterior probabilities of hypotheses are their prior plausibilities. Furthermore, if hypotheses are empirically distinct enough that evidence can drive the likelihood ratios to extremes, then subjective prior plausibility assessments become completely overridden by the evidence. This, at least, is what Bayesian inductivists hope for.

3. Likelihoods

The distinctive mark of a Bayesian support function is the logical or quasi-logical character of its likelihoods. That is, the likelihoods in Bayesian support functions are supposed to express what the hypothesis or theory *says* about the evidence. Bayesian logicians like Keynes and Carnap recognized this logical character, and it led them to the idea that inductive probability might be completely logicized, if only logic could be made to determine the values of prior probabilities as well. Keynes and Carnap each tried to implement this idea through syntactic versions of the principle of indifference. But logical form alone cannot determine reasonable values for prior probabilities, as examples employing Goodmanian grue-predicates illustrate. Still, the quasi-logical nature of the likelihoods is often apparent, especially when a hypothesis is deductively or statistically related to the evidence by direct inference likelihoods.

Even the most staunchly subjectivist Bayesians tend to rely on the special status of likelihoods. Consider, for instance, the following passages from the well-known Bayesian personalist paper by Edwards, Lindman, and Savage — the paper where they introduce the theory of stable estimation, a kind of Bayesian convergence result. Referring to a form of Bayes's Theorem, $P[H | D] = P[D | H] \cdot P[H] / P[D]$, as 'Equation 2', they say this:

In particular, all probabilities are conditional. Thus, $P[H]$ is the probability of the hypothesis H for you conditional on all that you know, or knew, about H prior to learning D ; and $P[H | D]$ is the probability of H conditional on that same background knowledge together with D .

Again, the four probabilities in Equation 2 are personal probabilities. ... But some are, so to speak, more personal than others. In many applications practically all concerned find themselves in substantial agreement with regard to $P[D | H]$, or $P[D | H]$ is public, as we say. This happens when $P[D | H]$ flows from some simple model that the scientist, or others, concerned accept as an approximate description of their opinion about the situation in which the datum was obtained.

... The other probabilities in Equation (2) are not at all public. Reasonable men may differ about them, even if they share a statistical model that specifies $P[D | H]$. People do, however, often differ much more about $P[H]$ and $P[D]$ than about $P[H | D]$, for evidence can bring initially divergent opinions into near agreement.

The probability $P[D]$ is usually of little direct interest, and intuition is often silent about it. It is

C_k there are possible outcomes O_{kv} on which the likelihoods differ: $P[O_{kv} | H_i \cdot B \cdot C_k] \neq P[O_{kv} | H_j \cdot B \cdot C_k]$. Otherwise H_i and H_j are empirically equivalent and no amount of evidence can support one over the other. If the true hypothesis has empirically equivalent rivals, then the convergence theorem shows that the odds against *the disjunction* of it with these rivals very probably goes to 0, and the posterior probability of *the disjunction* goes to 1. Among empirically equivalent hypotheses, the ratio of posterior probabilities equals the ratio of their priors: $P_a[H_j | B \cdot C^n \cdot E^n] / P_a[H_i | B \cdot C^n \cdot E^n] = P_a[H_j | B] / P_a[H_i | B]$. So the true hypothesis comes to have a posterior probability near 1 (upon driving the posteriors of empirically distinguishable rivals near 0) just in case non-evidential plausibility considerations make its prior plausibility much higher for the agent than the sum of the priors of its empirically equivalent rivals.

typically calculated or eliminated as follows. When there is a statistical model, H is usually regarded as one of a list, or partition, of mutually exclusive and exhaustive hypotheses H_i such that $P[D | H_i]$ are all equally public, or part of the statistical model. Since $\sum_i P[H_i | D] = 1$, Equation 2 implies that $P[D] = \sum_i P[D | H_i] P[H_i]$. (Edwards, Lindman, Savage, pp. 199-200)

Notice how much these Bayesians rely on the publicness and intuitiveness of likelihoods. Even the expectedness, $P[D]$, is to be calculated in terms of likelihoods, because the agent's intuitions are often *silent* about its value, whereas the likelihoods are better known — intuitions are far from silent about them. It is this tendency of likelihoods to be objective to a greater degree than other factors that gives Bayesian induction its distinctive character. Although subjectivist Bayesians often say they take probabilistic coherence (i.e. consistency with the standard probabilistic axioms) as the only constraint on a rational agent's belief strengths, mere probabilistic coherence cannot account for the central role Bayes's Theorem plays in the Bayesian account of hypothesis testing. Bayes's Theorem is central to Bayesians precisely because they treat the likelihoods as more objective than the posterior probabilities that likelihoods are used to calculate. Bayesians treat the likelihoods as stable points to which prior and posterior probabilities must coherently conform.

Where scientific hypotheses are concerned, the high degree of objectivity enjoyed by likelihoods is not just some fortunate accident. Recall the Popperian view that the demarcation between science and non-science rests with the falsifiability of scientific hypotheses, due to their logical entailment of evidential claims (Popper, 1959). This view is clearly too extreme. But there does seem to be something in it. The objectivity of science does seem to rely to a great extent on the objectivity of the connection between scientific hypotheses and the evidential claims that test them. For Bayesians, the likelihoods are supposed to capture this connection.

Imagine what a science would be like if the likelihoods were highly subjective. Its practitioners proceed by *interpreting* each theory to see what it says is likely to happen in specific circumstances. Although they carefully specify the auxiliaries and other background information they use, their interpretations of *what the theory says* about empirical matters are highly personal. So, in the assessment of the evidential support for competing hypotheses, it may well happen that likelihood ratios $P_\alpha[E^n | H_1 \cdot B \cdot C^n] / P_\alpha[E^n | H_2 \cdot B \cdot C^n]$ become quite large for scientist α , strongly supporting H_1 over H_2 , whereas the corresponding likelihood ratios $P_\beta[E^n | H_1 \cdot B \cdot C^n] / P_\beta[E^n | H_2 \cdot B \cdot C^n]$ for her colleague β become very small, refuting H_1 relative to H_2 . The very same evidence stream that supports H_1 and refutes H_2 for α does just the opposite for β . Under such circumstances theories do not *speak for themselves* about empirical matters. Rather, each practitioner is a sort of medium through which the theory speaks, and it says quite different things through each of its messengers. Although they employ the same *theoretical sentences*, each scientist understands the empirical import of these sentences so differently that each has in effect a distinct theory.⁷

My point is that the objectivity of the sciences rests largely on the abilities of scientific theories to *speak for themselves* about the evidence — to express largely the same empirical content to all members of a scientific community. The old deductivist/falsificationist ideal would have them do just that. Theories speak to the evidence via logical entailment; so there is no room for 'interpretative disagreement' over empirical content. But the deductivist ideal is clearly too strict. Presumably, hypotheses that bear on the evidence via *statistical direct inferences* may speak to the evidence in a somewhat similar, logically or quasi-logically objective manner. Even in cases where scientific theories are somewhat less formally rigorous, and so are tied to the data somewhat more loosely, the empirical objectivity of the science largely relies on a high degree of inter-subjective agreement about how likely evidence claims should be on various alternative hypotheses. It is primarily through such likelihoods that the empirical content of scientific theories is expressed.

In mathematically rigorous sciences an essential part of learning a theory involves acquiring the ability to logically deduce useful empirical consequences and statistical relationships. And even in less mathematically formal sciences, an essential part of learning a theory involves acquiring the ability to discern what the theory

⁷ This claim should not be confused with logical positivism. Positivists say '*same empirical import implies same theory*'. I am asserting the much less controversial claim, '*distinct empirical import implies (effectively) distinct theories*'.

says or implies about observable phenomena. Students are *trained up* on examples that instill a capacity to correctly perceive the implications. This training ultimately tends to provide a high degree of expert agreement regarding what a theory *says* about specific cases — i.e. a high degree of expert agreement on the values of the likelihoods, or at least on values of likelihood ratios.⁸ These agreed likelihoods do not simply represent an expert's individual degree-of-belief that an event has or will occur if a given theory is true — at least not in the first instance. Rather, it seems, experts consult the likelihoods they have internalized through their training to help them determine what to believe in specific cases, and how strongly to believe it.

Likelihoods are not conditional degrees-of-belief. They are much more like logical or analytic relationships.⁹ They represent what the hypothesis *says* or implies about the evidence, not how strongly the agent would believe the evidence if the hypothesis were added to everything else she knows. Degree-of-belief probabilities, even so-called counterfactual ones, are ill suited to this kind of objectivity or inter-subjective agreement required of the likelihoods. In the next section I will try to convince you of this. My point will be that there is important work for the *degree-of-belief* notion to do, but the *degree-of-support* notion associated with objective or public likelihoods does different work. Taking *degree-of-support* probabilities to be degrees of belief, even counterfactual ones, forces the *degree-of-belief* conception into a mold that doesn't suit it given the other work it does. Better to have two distinct notions and an account of how they interact.

4. Old evidence and related problems for the objectivity of degree-of-belief likelihoods

In typical Bayesian fashion, let Q_α represent an agent α 's rational *degree-of-belief* function. It is a probability function satisfying the usual axioms that represents the agent's present *belief-strengths* in various statements. Belief functions provide an idealized model of belief strength, and α is modeled as an ideally probabilistically rational agent. For example, Q_α assigns *degree-of-belief* 1 to all logical truths, so the agent is supposed to be logically omniscient and probabilistically coherent to an extent beyond the capabilities of real people. Belief functions extend the notion of ideally consistent belief. There is no harm in such idealization. We know how it is supposed to apply as a normative guide to real decision-making. An agent is supposed to attempt to make decisions based on her belief-strengths about the state of the world, her belief strengths about possible consequences of actions, and her assessment of the utilities of these consequences. This is all standard Bayesian fare.

4.1 The standard problem of old evidence applied to likelihoods

On the Bayesian model of belief, Q_α represents α 's net degree of belief in each statement. The agent's belief strength for a hypothesis, $Q_\alpha[H]$, is constrained by her belief strengths in other statements, including her belief strength in evidence statements, $Q_\alpha[E]$. Indeed, $Q_\alpha[H]$ incorporates the net impact of *all information* possessed by α on her belief strength in H . This is just what Q_α should do for the purposes of decision-making. For that purpose Q_α needs to represent α 's present belief strength for each statement that may be relevant to a decision. However, this tends to make Q_α ill-suited to represent the degree to which individual pieces of already known evidence *support* a hypothesis. For, if α is certain of some bit of evidence E , then $Q_\alpha[E]$ must be 1. In

⁸ Equations (2)-(4) of section 2 show that on the Bayesian account of confirmation the influence of the evidence is completely captured by *ratios of likelihoods*. So, although agreement on the *likelihoods* themselves is highly desirable, it will often suffice if we can get fairly objective or public values for *likelihood ratios*. Often in the softer sciences this may be the best one can do. The problems I will raise for the objectivity of *degree-of-belief likelihoods* apply equally to *degree-of-belief likelihood ratios*.

⁹ The standard axioms for Bayesian probability functions would have all conditional probabilities, including likelihoods, *defined* as ratios of unconditional probabilities. This may raise the worry that likelihoods must inherit whatever subjectivity there is in the unconditional probabilities that define them. To avoid such worries I think that degree-of-support functions are best understood as semantic conditional probability functions like those explicated by Field (1977). For more on this point see the discussion by Hajek (2003).

that case $Q_\alpha[H | E]$ must equal $Q_\alpha[H]$.¹⁰ Thus, the way in which evidence E supports H by increasing α 's belief strength in H does not overtly show up in Q_α . This is just the standard problem of old evidence.¹¹

Notice that when α is certain that E , the likelihood $Q_\alpha[E | H]$ must also be 1, even though H itself may say that E is quite unlikely. This wreaks havoc with the use of likelihoods and likelihood ratios in the assessment of the support of H by E . For, given any alternative hypothesis, the likelihood ratio $Q_\alpha[E | H^*] / Q_\alpha[E | H] = 1/1$. This undermines the point of drawing on Bayes's theorem to tell us how the evidence supports a hypothesis.

Consider an example. Suppose E states that the coin turns up *heads* 10 times in 10 tosses, and H says that the coin has a propensity $r < 1$ (e.g. $1/2$) to come up heads when tossed. And suppose it happens that E is true and the agent knows it. The likelihood we want to use in assessing the evidence for H is $P[E | H] = r$ ¹⁰, not $Q_\alpha[E | H] = 1$. Where *certain* evidence is concerned, the likelihoods and likelihood ratios expressed by Q_α utterly fail to capture the way in which *what hypotheses say* about the evidence is supposed to bring about their evidential refutation or support. This difficulty for Bayesian confirmation theory is part and parcel of the standard problem of old evidence.

The standard old evidence problem applies only to evidence that the agent knows with certainty. Some Bayesians may hope to live with this problem, confident that for all evidence not yet known with certainty Q_α still captures correct likelihoods and permits Bayes's theorem to express the appropriate confirmation relation. However, the standard problem of old evidence is only the tip of an iceberg. It turns out that the incorporation into Q_α of even very weak information 'involving' possible evidence E may completely undermine the public or objective values of likelihoods expressed by Q_α . This should be no surprise. The agent's belief function likelihoods $Q_\alpha[E | H]$ reflect her *total degree-of-belief* in E , based on H together with *everything else* she believes about E . So the agent's present belief function may capture appropriate, public likelihoods for E only if E is completely isolated from the agent's other beliefs. And this will rarely be the case. To see the point, let's look at two kinds of information about E that fails to make E certain, but nevertheless undermines the ability of belief functions to capture appropriate evidential likelihoods.

4.2 The problem of old disjunctively-entangled evidence

Suppose the agent isn't at all certain of E , and her present belief-function likelihood is $Q_\alpha[E | H] = r < 1$. And suppose she is also less-than-certain about whether some disjunctive claim ($E \vee D$) will hold if H does — i.e. $Q_\alpha[E \vee D | H] = s < 1$. If the agent then becomes certain that ($E \vee D$) (and if this is her strongest newly certain belief), then her belief function should be updated in the usual Bayesian way. Her new belief function $Q_{\alpha\text{-new}}$ is related to her previous belief function by the formula: $Q_{\alpha\text{-new}}[S] = Q_\alpha[S | E \vee D]$. Thus, $Q_{\alpha\text{-new}}[E | H] = Q_\alpha[E | H \cdot (E \vee D)] = Q_\alpha[(E \vee D) \cdot E | H] / Q_\alpha[E \vee D | H] = Q_\alpha[E | H] / Q_\alpha[E \vee D | H] = r/s > r$. So, where the old likelihood was $Q_\alpha[E | H] = r$, the updated likelihood must become $Q_{\alpha\text{-new}}[E | H] = r/s$.

This kind of belief updating is devastating to the objectivity of the likelihoods. For example, let H say that a specific coin is fair and E say that a specific toss comes up *heads*. And let us suppose that α 's belief function likelihood is the usual objective likelihood for heads on tosses of a fair coin, $Q_\alpha[E | H] = 1/2$. Now, α learns from a friend who has been tossing the coin recently that either the coin has come up *heads* on the specific toss at issue *or* it came up *heads* on the three previous tosses — i.e., α learns ($E \vee D$), where D states that the past three tosses were *heads*. And let's suppose α takes the outcomes of each toss to be independent given H . Then $Q_\alpha[E \vee D | H] = Q_\alpha[E | H] + Q_\alpha[D | H] - Q_\alpha[E | H] \cdot Q_\alpha[D | H] = (1/2) + (1/8) - (1/2) \cdot (1/8) = 9/16$. Thus, $Q_{\alpha\text{-new}}[E | H] = (1/2)/(9/16) = 8/9$. That is, α 's degree-of-belief likelihood for *heads* on the toss of the fair coin, given everything else she now know, has become $8/9$.

Don't get me wrong. Given what she knows, $Q_{\alpha\text{-new}}[E | H] = 8/9$ may be exactly the right betting quotient on *heads* for α . But this is clearly not the kind of objective or public likelihood that Bayesians want to draw on when using E as evidence for or against the *fairness hypothesis H*.

¹⁰ It is usual in such discussions to suppress B and C ; perhaps they are supposed to have degree-of-belief 1, or very nearly 1. For the sake of simplicity I will follow this convention for now.

¹¹ The problem of old evidence was first raised by Glymour in his, 'Why I Am Not a Bayesian,' (1980, Ch. 3). See Eells (1985) and Earman (1992, Ch. 5) for very thorough analyses of the problem.

The problem is *not* that using this likelihood would necessarily give the wrong confirmational result. Indeed, provided she started out with the correct objective likelihood, and then updated it as in the example, she may get the correct confirmational result from her new likelihood together with her updated ‘prior’ probability $Q_{\alpha\text{-new}}[H]$.¹² Rather, the problem is that if such personal belief-function likelihoods were the only likelihoods a Bayesian could draw on, and if we had to use these in scientific confirmational contexts, then each member of the scientific community would generally have his or her own distinct personal likelihood for the evidence, depending on what else he or she knows about the evidence. This would make a complete hash of scientific hypothesis testing. What should a researcher report in a scientific journal? That, given what *she* knows, *she* takes the likelihood of the experimental outcome E on H to be r , but that *you* should consult everything that *you* know to come up with your own likelihood value for E on H ? Presumably not. Rather, in discussing the evidence she should refer to public, support function likelihoods, not personal degree-of-belief likelihoods. I’ll return to this point in the next section, after we first see that things are even worse than this example suggests.

To see how troublesome this effect can be, let us consider one further example. As part of a periodic physical exam a physician intends to have her patient tested for hepatitis with a standard blood test. She knows from medical studies that there is a 5% false positive rate for this test; so her belief strength that the test result will be positive, E , given hepatitis is not present is $Q_{\alpha}[E | \neg H] = .05$. Now, the lab is normally very reliable. So, if asked, the physician would say her belief strength that the lab would mix up the blood work for her patients, D , is very small. Indeed, given that her patient is free of the disease, she takes the likelihood of a mix-up to be around .01 (i.e. $Q_{\alpha}[D | \neg H] = .01$), and she finds it equally likely regardless of whether her patient’s test is positive (i.e. $Q_{\alpha}[D | \neg H \cdot E] = .01$). Then the physician receives a call from the lab. The technician, seems unsure of the reliability of the result, and says, ‘your patient came up positive or we’ve mixed up his results. Do you want to take another blood sample and have us run the test again?’ The physician’s new belief-function-likelihood for a false positive on her patient becomes $Q_{\alpha\text{-new}}[E | \neg H] = Q_{\alpha}[E | \neg H \cdot (E \vee D)] = Q_{\alpha}[E | \neg H] / Q_{\alpha}[E \vee D | \neg H] = Q_{\alpha}[E | \neg H] / (Q_{\alpha}[E | \neg H] + Q_{\alpha}[D | \neg H] - Q_{\alpha}[D | \neg H \cdot E] \cdot Q_{\alpha}[E | \neg H]) = .05 / (.05 + .01 - (.05)(.01)) = 1 / 1.19 \approx .84$.

Indeed, even rather trivial knowledge of disjunctive claims involving E may, when learned with certainty, completely upset the likelihoods for an agent’s belief function. And an agent will almost always have some such trivial knowledge. For example, the physician may also learn such new disjunctive facts involving E as, ‘the patient’s insurance company won’t cover additional tests and treatment *unless* the test turns out positive’ and ‘either the test is positive *or* it turns out that thirty-seven of my patients will have had a negative hepatitis test this year’, etc. Updating on such information can force a physician’s *belief function* to deviate widely from the evidentially relevant objective, textbook values of test result likelihoods. This is not just the usual problem of old evidence; let’s call it the *problem of old disjunctively-entangled evidence*.

4.3 The problem of old uncertain evidence

Consider the following situation. A physician’s objective, textbook values for the likelihood of a positive sputum test, O , if the patient has lung cancer, H , is $Q_{\alpha}[O | H] = .9$; and her objective likelihood for a positive test when no cancer is present is $Q_{\alpha}[O | \neg H] = .01$. Based on this particular patient’s risk factors, she assigns a prior probability for lung cancer of $Q_{\alpha}[H] = .1$. While examining the patient, she notices his persistent cough. Based on the cough and other clinical impressions, she fears that the sputum test may well come back positive. On this basis her ‘intuitive’ new belief strength for a positive test becomes $Q_{\alpha\text{-new}}[O] = .5$ — this is how strongly she has now come to believe that the test will be positive, based on her examination. However, she recognizes the result of a sputum test to be definitive enough to

¹² For example, if testing H against alternative H^* , it will turn out that $Q_{\alpha\text{-new}}[H^* | E] / Q_{\alpha\text{-new}}[H | E] = Q_{\alpha}[H^* | E] / Q_{\alpha}[H | E]$, since $Q_{\alpha\text{-new}}[H^* | E] / Q_{\alpha\text{-new}}[H | E] = (Q_{\alpha\text{-new}}[E | H^*] / Q_{\alpha\text{-new}}[E | H]) \cdot Q_{\alpha\text{-new}}[H^*] / Q_{\alpha\text{-new}}[H] = [(Q_{\alpha}[E | H^*] / Q_{\alpha}[E \vee D | H^*]) / (Q_{\alpha}[E | H] / Q_{\alpha}[E \vee D | H])] \cdot Q_{\alpha}[H^*] / Q_{\alpha}[H] = (Q_{\alpha}[E | H^*] / Q_{\alpha}[E | H]) \cdot (Q_{\alpha}[H^*] / Q_{\alpha}[H]) = Q_{\alpha}[H^* | E] / Q_{\alpha}[H | E]$.

*screen off*¹³ the cancer hypothesis from the observed cough and other clinical impressions: $Q_{\alpha\text{-new}}[H | O] = Q_{\alpha}[H | O]$ and $Q_{\alpha\text{-new}}[H | \neg O] = Q_{\alpha}[H | \neg O]$. It turns out that under these circumstances, as she sends the sample off to the lab, her new belief-function likelihood *must* become $Q_{\alpha\text{-new}}[O | H] = .988$ for a true positive test — significantly higher than the .9 textbook value. Similarly, her new false positive rate *must* become $Q_{\alpha\text{-new}}[O | \neg H] = .084$ — more than eight times higher than the textbook rate of .01.¹⁴ This produces a significant change in her likelihood ratios as well, from $Q_{\alpha}[O | H]/Q_{\alpha}[O | \neg H] = 90$ to the much lower value $Q_{\alpha\text{-new}}[O | H]/Q_{\alpha\text{-new}}[O | \neg H] = 11.76$. Let's call this problem, where updating the belief strengths of possible evidential outcomes wrecks the public values of degree-of-belief likelihoods, the *problem of old uncertain evidence*.

The general form of this problem is this. The agent α is faced with a partition (i.e. a mutually exclusive and exhaustive set) of alternative hypotheses $\{H_k\}$. Among the evidence claims that may bear on these hypotheses are the alternative outcomes in a partition $\{O_i\}$. At a time before the relevant observation is made the agent's present degree-of-belief function is Q_{α} . Let's suppose for now that the belief function likelihoods $Q_{\alpha}[O_i | H_k]$ have proper public or objective values. The agent then learns some 'new bit of information' that changes her degrees of belief (from $Q_{\alpha}[O_i]$ to $Q_{\alpha\text{-new}}[O_i]$) regarding which of the outcomes O_i will be found when the appropriate observation is eventually made. (This new information may involve 'hard data' relevant to the possible outcomes O_i , or it may just be a hunch; for our purposes it doesn't matter.) Furthermore, suppose this new information does not affect what the agent's belief strengths for hypotheses would become if some one of the O_i were found to be true. Then, in almost all cases updating *must* cause at least some of the agent's likelihoods to deviate from their public values.

More formally, what I'm claiming is this. Suppose that the new information changes the agent's degree-of-belief function from Q_{α} to $Q_{\alpha\text{-new}}$, updating her belief strengths in at least some of the possible outcomes, $Q_{\alpha\text{-new}}[O_i] \neq Q_{\alpha}[O_i]$; and also suppose that the conditional belief strengths for the H_k on the various O_i are maintained, $Q_{\alpha\text{-new}}[H_k | O_i] = Q_{\alpha}[H_k | O_i]$.¹⁵ It turns out that in the presence of this kind of belief updating, whenever a very plausible *additional condition* holds,¹⁶ some of the agent's likelihoods must change — at least some of the new likelihoods $Q_{\alpha\text{-new}}[O_i | H_k]$ must differ from the corresponding old likelihoods $Q_{\alpha}[O_i | H_k]$. If her original likelihoods were public, at least some of them must turn private after updating. This means that once α updates to $Q_{\alpha\text{-new}}$ she can no longer account for her present belief strength for H_k given O_i in terms of Bayes's theorem together with the *public* or *objective* likelihoods.

How bad is this for Bayesians? After all, although the agent's likelihoods are no longer public, as stated this problem takes her posterior-probability-belief-strengths for hypotheses, $Q_{\alpha\text{-new}}[H_k | O_i] = Q_{\alpha}[H_k | O_i]$, to

¹³ To say that O_i *screens off* H_k from propositionally expressible information E just means that H_k is probabilistically independent of E given O_i — i.e. $Q_{\alpha}[H_k | O_i \cdot E] = Q_{\alpha}[H_k | O_i]$. Intuitively the idea is that the information content of O_i makes the addition of E irrelevant to H_k . The same idea applies to cases of Jeffery Updating, where the information that brings about an update from Q_{α} to $Q_{\alpha\text{-new}}$ is not propositionally expressed, but maintains the relationship $Q_{\alpha\text{-new}}[H_k | O_i] = Q_{\alpha}[H_k | O_i]$. In this case we simply say that O_i *screens off* H_k from the *new-update information*.

¹⁴ To calculate $Q_{\alpha\text{-new}}[O|H]$ first notice that $Q_{\alpha}[O] = Q_{\alpha}[O|H] \cdot Q_{\alpha}[H] + Q_{\alpha}[O|\neg H] \cdot Q_{\alpha}[\neg H] = .099$. Now let $R = Q_{\alpha\text{-new}}[\neg O|H]/Q_{\alpha\text{-new}}[O|H] = (Q_{\alpha\text{-new}}[H|\neg O] \cdot Q_{\alpha\text{-new}}[\neg O]) / (Q_{\alpha\text{-new}}[H|O] \cdot Q_{\alpha\text{-new}}[O]) = (Q_{\alpha}[H|\neg O]/Q_{\alpha}[H|O]) \cdot (Q_{\alpha\text{-new}}[\neg O]/Q_{\alpha\text{-new}}[O]) = Q_{\alpha}[\neg O|H]/Q_{\alpha}[O|H] \cdot (Q_{\alpha}[O]/Q_{\alpha}[\neg O]) \cdot (Q_{\alpha\text{-new}}[\neg O]/Q_{\alpha\text{-new}}[O]) = 11/901$. Then $Q_{\alpha\text{-new}}[O|H] = 1/(1+R) = 901/912 = .988$. $Q_{\alpha\text{-new}}[O|\neg H]$ may be calculated similarly.

¹⁵ This is the condition that Richard Jeffrey calls *rigidity* (1965, 1970, 1988, 1991). It is an essential feature of his well-known account of updating on uncertain evidence. In (Hawthorne, 2004a) I discuss *rigidity* in more detail, and extend Jeffrey-updating to sequences of updates in a way that makes it independent of update order.

¹⁶ Each of several plausible *additional conditions* would suffice. One such condition is this: at least one of the hypotheses assigns a non-zero likelihood to each outcome — i.e. an H_k makes $Q_{\alpha}[O_i | H_k] > 0$ for all O_i . A different sufficient condition is this: some outcome gets a non-zero likelihood from each hypothesis — i.e. an O_i makes $Q_{\alpha}[O_i | H_k] > 0$ for all H_k . These two conditions are each special cases of a much weaker condition, which is spelled out in the Appendix. The Appendix also contains a proof of this result.

retain the values they had for the earlier belief function, where the likelihoods were (assumed to be) public. The problem is that in general the agent's original belief-strength likelihoods might well not be public. The point of the above discussion is to show that *even if they were objective or public*, updating would change the likelihoods, destroying whatever objectivity they had. Furthermore, even in cases where an earlier belief function had public likelihoods, additional updates on other evidence will usually change the values of the posterior-probability-belief-strengths. (That's the point of gathering evidence, isn't it?) After that happens, the posterior-probability-belief-strengths lose all apparent relationship to any earlier public belief-strength-likelihoods there may have been.¹⁷

Consider a further effect of this kind of uncertain updating of evidence claims. In a given confirmational context, due to their individual experiences, different agents each have somewhat different partial evidence for evidence claims O_i ; so, each updates his belief strengths in the O_i differently. Thus, even if these agents start out with the same likelihoods for the possible evidence claims, the values of their *current degree-of-belief likelihoods* will come to disagree. Then, when outcome O_j is finally 'observed' and the agents all become certain of it, each agent must employ his own personal *current belief function likelihood* for O_j to update his belief strengths for the hypotheses. This makes a hash of the usual view of Bayesian confirmation, where likelihoods are supposed to be public, and where the only significant disagreement among agents is supposed to be over prior probabilities.

4.4 Counterfactual belief-strengths?

Section 4.1 described the standard problem of old evidence. The most prominent Bayesian response to this problem is that to see the influence of old evidence on a hypothesis we should employ some alternative degree-of-belief function — not the agent's current belief function. As Howson and Urbach explain it, we should employ a counterfactual degrees-of-belief function that gauges the support of H by E , '...according to the effect which one believes a knowledge of E would now have on one's degree of belief in H , on the (counter-factual) supposition that one does not yet know E .' (1993, pp. 404-5)

Howson and Urbach acknowledge that there will often be no uniquely correct way to subtract E from the agent's current stock of knowledge; so there is no uniquely correct counterfactual degree-of-belief function that expresses the proper effect of E on H . But, they claim, there will often be a 'satisfactory' way to subtract E from current knowledge. The idea is that in realistic cases a person's total body of knowledge is the set of logical consequences of some set of sentences K , which are, in effect, the person's axioms for her body of knowledge. And, they contend, E will often be an independent axiom of K . When this is the case, $K - \{E\}$ picks out the revised stock of knowledge uniquely. But, they admit, in cases where E is not an independent item in K , '...we may just have to conclude that what the probability of E , and of E given H , would be, were one counter-factually assumed not to know E , is not defined. But that such counter-factually based probabilities do not always exist should not blind us to the fact that in many cases they do, and consequently that in many cases support relative to known evidence is also perfectly meaningful.' (1993, p. 406)

Howson and Urbach's attempt is valiant, and I would welcome it as an explication of degree-of-support functions, if it worked. But I don't think that it can work. With regard to the *standard problem of old evidence*, even if E is an independent member of K , simply removing it doesn't tell us what the revised, counterfactual belief function CFQ_α should be. Whereas $Q_\alpha[E] = 1$, $CFQ_\alpha[E]$ should be less than 1; but what should its value be? And what value should $CFQ_\alpha[H]$ and $CFQ_\alpha[E|H]$ have if we remove E from K ? Howson and Urbach give

¹⁷ That is, after updating the belief strengths of the O_i by first updating from Q_α (whose likelihoods are public) to $Q_{\alpha\text{-new}}$ (whose likelihoods are not public), the values of $Q_{\alpha\text{-new}}[H_k | O_i]$ still agreed with the values of $Q_\alpha[H_k | O_i]$, which are the values one would have gotten from the public likelihoods. But additional updating on other, unrelated evidence will usually make the most current belief function, $Q_{\alpha\text{-newest}}$ exhibit a quite different posterior probability — e.g. on learning that D , $Q_{\alpha\text{-newest}}[H_k | O_i] = Q_{\alpha\text{-new}}[H_k | O_i \cdot D] \neq Q_{\alpha\text{-new}}[H_k | O_i] = Q_\alpha[H_k | O_i]$. So, at that point the agent's current belief function likelihoods are no longer public, and her posterior probabilities for her current belief function, $Q_{\alpha\text{-newest}}[H_k | O_i]$, no longer have the same values she would have gotten from the belief function Q_α by employing public likelihoods.

us no guidance on this. I suppose the idea is that once the agent moves to a stock of knowledge where the degree of belief in E is less than 1, everything else will fall into place. Her counterfactual-belief-likelihoods $CFQ_{\alpha}[E | H]$ will automatically assume appropriate public values.

In response to this proposal I first want to point out that in many cases we seem able to assess the values of likelihoods perfectly well without considering whether E is an independent axiom of our stock of knowledge. And it seems to me that in many cases where we do know the appropriate likelihoods, E is not an independent axiom of our current stock of certain knowledge. So it seems rather implausible that our likelihoods are gotten by removing E from a stock of knowledge, and then assessing the resulting belief strengths.

Worse yet, the problems of *old disjunctively-entangled evidence* and of *old uncertain evidence* show that simply making $CFQ_{\alpha}[E]$ less than 1 (e.g. by removing E from K) will not generally leave the likelihoods $CFQ_{\alpha}[E | H]$ sufficiently unconstrained to permit them to assume the appropriate public values. Section 4.2 shows that if the agent is certain of almost any contingent claim of form ' $(E \vee D)$ ', and if this claim remains certain for her counterfactual belief function, then the counterfactual likelihood cannot take on the public value it might otherwise have. (And contingent claims of form ' $(\neg E \vee D)$ ' cause similar problems.) So, if the counterfactual belief function idea is to work, then, at the very least, the agent must remove all disjunctive sentences (and material-conditional sentences) involving E from her stock of knowledge as well. But, as Howson and Urbach point out in the case of removing E , this may work only if *all* such conditional claims occur as 'independent axioms' in K . It seems doubtful that this is often the case.

If the problem with conditional entanglements isn't bad enough, the *problem of old uncertain evidence* makes matters almost impossible for the counterfactual belief function idea. If the counterfactual likelihood, $CFQ_{\alpha}[E | H]$ is to take an appropriate public value, then the counterfactual probability must be based on a stock of knowledge from which *all* information relevant to E has been removed, and this must be done in a way that leaves the rest of the stock of knowledge intact. It would be difficult enough to do this for all propositionally expressible knowledge D that bears on E . But we must also remove any non-propositional knowledge or experiential influences that make the degree of belief in E different than it would be in the absence of such knowledge or influences. How are we to do that?

It looks like we are left only with the following strategy: if the agent knows the objective or public values that likelihoods should have, then she is to select a counterfactual belief function that gets the likelihoods right. That is, we simply legislate that the counterfactual likelihoods $CFQ_{\alpha}[E | H]$ are to take on the public values we know they should have. Perhaps the agent is to consider the closest world in which her likelihoods have the right values, and her counterfactual belief function should then assign probability 1 to all of the knowledge she would still have at that world. But in order to take advantage of this strategy the agent must already have the objective or public likelihoods at her disposal, independently of the counterfactual belief function she is supposed to construct. And, to the extent that she already has access to the public likelihoods, she has the essential part of a degree-of-support function already at hand.

The idea that likelihoods are counterfactual belief strengths looks by now like quite a stretch. And it seems psychologically unrealistic as well. We don't seem to actually determine likelihoods by stripping away interfering information from our body of knowledge, adding H to whatever remains, and then figuring out how much confidence we would have in E . So, public or objective likelihoods are not based on counterfactual belief strengths either in principle or in fact. But we do often seem to have objective or public likelihoods available to us, especially in the sciences. And the empirical objectivity of the sciences apparently demands them.

5. Why Bayesian confirmation requires both functions

I think that support functions are best understood as a kind of logical or quasi-logical probability. But unlike the kind of logical function championed by Keynes and Carnap, it does not depend only on the logical structure of sentences. Clearly logical form alone cannot dictate appropriate probabilities of all statements. Rather, the likelihoods seem to be logical or quasi-logical in a sense somewhat analogous to the sense in which the 'analytic entailment' of some sentences by others is *logical*. Likelihoods represent the empirical content of hypotheses — what the hypotheses *say* about the evidence — rather than what an agent is supposed to believe when the hypothesis is added to everything else she believes.

However, this interpretational issue is not the most important point I'm trying to make. My central contention is that a proper Bayesian model of scientific inference should represent agents as possessing two distinct probability functions. The first function represents the agent's total doxastic state — his *current belief-strengths* for statements, based on everything he knows. This is the function that Bayesians associate with decision-making. The main goal of scientific inquiry is to update this function, with the aim of ultimately producing high belief strengths for true hypotheses and low belief strengths for false ones. I maintain that in scientific contexts the updating of this function is mediated by a second, more static function that represents objective or public likelihoods. In this section I want to further press the argument for this two-function model. I hope to convince you of this, regardless of whether you buy my conception of the second function as *logical* or *quasi-logical*. So, for present purposes I want to disentangle the primary point about the need for the two-function model from my view about the *logical* nature of support functions. Thus, for now you may take a support function to be some kind of 'second belief function' — e.g. a 'counterfactual' belief function — if you wish.

My main contention is that in order to adequately model how belief strengths for hypotheses are influenced by evidence in scientific confirmational contexts, Bayesians need to supplement the usual *current-belief-strength functions* with a second probability function that maintains objective or public (perhaps counterfactual) likelihoods. I will spell out the formal details of how, on this two-function model, support functions should inform belief functions in the next section. But here I will further press the point that an adequate Bayesian model of how belief is altered by evidence in scientific contexts requires a role for support functions. Indeed, it would be nearly impossible for science to proceed as it does without drawing on something very much like support function likelihoods.

5.1 Modeling scientific practice

When new evidence becomes available, real scientific agents (regardless of whether they are Bayesians) do seem to draw on objective or public likelihoods, shared by the scientific community, to assess how strongly they now believe hypotheses based on that evidence. But, I have argued, these public likelihoods cannot normally come from an agent's present or past *actual belief function* (if he has such a thing), because those likelihoods are plagued by old evidence problems. So the likelihoods on which hypotheses are evaluated *must* be carried by a second function, which I call a *support function*.

A traditional subjectivist Bayesian might respond to this point as follows:

Sure, real scientists tend to appeal to public likelihoods. But there is no very good reason for scientific Bayesianism to model this particular facet of scientific practice. *Standard Bayesian kinematics* — i.e. the process of updating on new information D via the algorithm $Q_{\alpha\text{-new}}[S] = Q_{\alpha}[S | D]$, for all sentences S — shows how successive updates may occur through direct transformations from one *current belief function* to the next, with no need to appeal to a second, supporting probability function. The conditional probability $Q_{\alpha}[S | D]$ on which the update of α 's belief strengths are based need only employ α 's current belief function likelihoods; and that's all α needs in order to perform her Bayesian update, regardless of whether these likelihoods are objective or public. So the Bayesian approach doesn't really *need* this second function to account for the updating of belief. Why not stick with *standard kinematics*? Why bring this second probability function into Bayesian belief updating?

I want to say several things in answer to this point. First, don't get me wrong. I'm not claiming that the standard updating kinematics necessarily gives the wrong values for *updated current belief functions*. The kinematics of updating via support functions that I'm proposing should (under normal circumstances) result in precisely the same updated *current belief functions* as standard kinematics, but through a different, more realistic algorithm.¹⁸ Indeed, standard kinematics is in some ways more general

¹⁸ Isn't that exactly how 'counterfactual' probabilities are supposed to behave? Isn't the application of an

than support function kinematics in that the standard approach applies to updates of individual belief functions even in cases where there are no underlying (counterfactual) support functions to tie together the belief functions of diverse agents.¹⁹ The requirement that the agents in a scientific community update their belief functions in accord with their underlying support functions places additional constraints on the agents' belief functions. That is, much as the axioms of probability place coherence constraints on an individual agent's rational belief strengths, the requirement that the belief function of each agent in a given scientific community is tied to an underlying support function that shares likelihoods with the support functions of other agents places additional coherence constraints on his rational belief strengths.

What do I mean by the claim that updating mediated by support functions employs a 'more realistic' updating algorithm than standard Bayesian kinematics? Suppose scientists were to actually update belief strengths for hypotheses by following the standard kinematics process. The only likelihoods they would employ in belief updating would be their *current belief function likelihoods*. And, as we've seen, these will seldom maintain objective or public values. So, when various scientists disagree about the likelihoods hypotheses impart to evidence claims, the proper attitude is simply, 'that's your likelihood value, and here's mine.' That is, if the usual story were right, the only relevant discussion among scientists about the values of likelihoods would concern whether each agent's likelihood values fit coherently with the other personal belief strengths he has at the moment, and whether his updates from previous belief functions have been carried out properly. But this is not how real scientists behave, nor is it how they should behave. Rather, when scientists disagree about the likelihood values imparted by hypotheses, they usually try to convince one another about what the *appropriate* (counterfactual) values should be. They do so because they take getting the likelihoods right, and public, to be central to the way evidence is used to update beliefs about hypotheses in the sciences.

If scientists actually updated their belief-strengths for hypotheses according to the standard kinematics story, imagine how the discussion might go at a scientific conference where a new experimental or observational result O is first announced. You might overhear a snippet of α 's conversation with her colleague β that runs like this:

Great news that they've observed O ! This result changes my confidence in H_1 quite a bit, since *my* most recent likelihood for O on H_1 was $Q_\alpha[O | H_1] = .9$ while *my* most recent likelihood for this result on H_1 's strongest competitor, H_2 , was only $Q_\alpha[O | H_2] = .2$. How about you, β ? How likely was O according to H_1 for *you*, and how likely did *you* take O to be on H_2 ? Were your likelihoods for these two hypotheses also different enough for O to count as significant evidence for H_1 over H_2 ? ... Or, was O perhaps more likely on H_2 than on H_1 for you?

That is, when it comes to the updating of belief-strengths for hypotheses, standard Bayesian kinematics can only draw on consideration of *your* likelihood and *my* likelihood. But in reality scientists often use common, publicly agreed (counterfactual?) likelihoods in the updating of their belief strengths. Indeed, if each scientist had only his own personal likelihoods, and they tended to disagree widely with the likelihoods of others, conducting empirical research and publishing it would hardly seem worth the investment of time and resources.

Standard kinematics is not even close to how real scientific agents update belief strengths for hypotheses; and there is no good reason to think that they should proceed that way. Rather, real agents update by appealing directly to support function likelihoods, or something very much like them. It should

agent's 'counterfactual' probability function to all evidence used in the sequence of *standard updates* supposed to result in the same *current belief function* he arrived at through the *standard update* sequence?

¹⁹ However, *support function kinematics* does have some advantages. For instance, *standard updating* fails for cases where the agent changes his mind about the evidence and wants to retract some previous update. It can be shown that *current belief functions* have no 'memory' of the previous belief functions from which they came — so a one-function model doesn't permit retractions. But the support-function update model handles retractions easily.

be quite apparent that this is at least implicitly how its done, since: (1) the public likelihoods real agents often use cannot (on pain of the old evidence problems) generally be *current belief function* likelihoods; and (2) agents in scientific confirmational contexts almost never compute *current belief function likelihoods*, much less *use them* to update their current belief functions on additional evidence. So, while I concede that the *usual* Bayesian kinematics story about getting each successive *current belief function* from the previous *current belief function* could be carried out in principle, I contend that this story is far removed from what real agents actually do, or should do, when updating their beliefs regarding scientific hypotheses. Thus, the two-function model of Bayesian belief formation, which draws on support-functions, is well worth pursuing in more detail.

5.2 Bayesian networks

Before turning to a discussion of the formal workings of how support functions should inform belief, one more point is worth thinking about. Let's consider how automated Bayesian networks are employed by real agents to determine appropriate belief strengths for hypotheses. Any automated Bayesian network, say a network designed to aid physicians in diagnosing patients on the basis of symptoms and test results, is essentially a network of objective or public ('counterfactual', if you insist) likelihoods. Physicians who use such a system accept the likelihoods built into it as having correct, publicly agreed values. Why else would a whole medical community be willing to employ the same network system (together with each patient's symptoms and test results) to assess appropriate belief strengths for disease hypotheses about their patients?

Notice that the support function built into a Bayesian diagnostic system cannot (on pain of old evidence problems) itself be intended to represent the physician's *current belief function* about her patient. Rather, the built-in support function is a logical tool that aids the agent in fleshing out her personal belief-strengths. Based on her own observation of the patient and on reports from lab tests, the physician comes to believe evidence claims about the patient's symptoms and test results. On a Bayesian model, the physician's belief-strengths for evidential claims are part of *her current belief function*. (In general the physician need not be *certain* of the symptoms and test results. Well-designed Bayesian network systems should have a built-in version of *Jeffrey updating* to handle *uncertain evidence*. See Jeffrey's (1991) and Hawthorne (2004a) for more details.) However, it is not easy for the physician to ascertain what her current belief strengths should be for all of the various possible diagnostic hypotheses about her patient. To do so she would need to internalize the textbook values of likelihoods connecting the various possible diseases to symptoms and test results, and then perform a lot of calculation. That's where the automated system comes in. She uses its support function, with its built-in textbook likelihoods, which she accepts as her own support function, to help her ascertain what her current belief strengths for disease hypotheses should be.

An automated Bayesian network typically operates by maintaining two distinct probability functions. It maintains the original network support function, the one with the public likelihoods, and it maintains a *current belief function*, which represents the belief strengths the system recommends as appropriate on the totality of the evidence it has been given. When updating on new evidence, such systems never overwrite their original support function. Rather, the system continues to employ likelihoods from the original support function to perform each new update on additional evidence. Indeed, such systems almost never compute *current belief function likelihoods* — i.e., they don't compute the likelihoods evidence claims would have if a disease hypothesis were added to everything else *currently believed* by the system with certainty. They could do so; but users seldom care about the values of *these* likelihoods. Such likelihoods are neither particularly useful to the agent, nor are they needed for use by the system. Rather, with the right algorithm, the original support function likelihoods suffice for all updating; and they turn out to be absolutely essential to updating in cases where the agent changes her mind about some bit of evidence, and wishes to retract a previous update.

Automated Bayesian networks generally have built-in prior probabilities for the various disease hypotheses, which represent the *base rates* of diseases in the general population. But before evidential updating begins, the system may permit the physician to override its built-in priors. She may input her

own prior probabilities for disease hypotheses — those she considers appropriate to her particular patient, based on information about his risk group. The physician then inputs her current belief strengths (usually certainties) about the presences of symptoms and test results. The system responds by computing updated *current belief strengths* of its own for the various disease hypotheses, based on its objective likelihoods. When the physician has entered into the system all of the relevant evidence she knows about her patient, she may then adopt the system's updated belief strengths as her own.

Bayesian subjectivists often treat counterfactual belief strengths as though they are only useful retrospectively, to show how some pieces of old evidence may count as support for a given hypothesis. On the standard model of Bayesian kinematics this is all that counterfactual belief functions are good for. But that model does not accurately reflect how real agents who employ automated Bayesian networks proceed. And it does not reflect how real agents proceed in scientific confirmational contexts more generally. In many contexts the scientific community is in wide agreement about what hypotheses or theories *say* or imply about various evidence claims, and thus about the appropriate values for support function likelihoods. When such likelihoods are available, scientists generally appeal to them to update their belief functions. Thus, if *support functions* are *counterfactual belief functions*, then counterfactual belief should play a central role in the Bayesian model of how evidence informs scientific belief.

6. The belief-support connection

Let us suppose that each member of a community of Bayesian agents possesses both a belief function and a support function. Presumably the degree to which the evidence supports a hypothesis, as represented by the agent's support function, should govern her belief strengths for hypothesis. Exactly how is this supposed to work? I will address this issue directly in a moment. But I think the answer will make better sense if we first attend to the nature of the agents we are discussing.

Ideal Bayesian agents are components of a formal model of probabilistic inference and decision. In such models we describe these agents in terms of their utility functions, belief functions, support functions, etc., and in terms of how these functions transform through updating on new information. Like all formal logics, the logic of such Bayesian models is normative. Bayesian logicians try to show that the ideal agents of their models gain certain benefits (e.g. they avoid *Dutch book*), and suggest that real agents may gain similar benefits through emulating the logic of their ideal brethren.

The agents in the following discussion are such Bayesian idealizations. I will proceed by constructing a formal model of scientific inference and belief that shows how support functions should inform the belief strengths of ideal Bayesian agents. Then I'll point out properties of this model that make it worthy of emulation by real agents.

6.1 Support functions and empirical contexts

Agents seldom share a support function. For they often disagree over the prior plausibilities of hypotheses. But all who understand the empirical import of hypotheses in the same way will possess identical (or nearly identical) support function *likelihoods*. More specifically, let us say that any two agents α and β *share an empirical context* for hypothesis partition $\{H_1, H_2, \dots\}$, background B , and observation conditions C_1, \dots, C_n, \dots , with respective outcome partitions $\{O_{1u}\}, \dots, \{O_{nv}\}, \dots$, *just in case* for each possible sequence of outcomes from the partitions, they agree on the values of support function likelihoods — i.e., $P_\alpha[O_{1u}\dots O_{nv} | H_i \cdot B \cdot C^n] = P_\beta[O_{1u}\dots O_{nv} | H_i \cdot B \cdot C^n] = P[O_{1u}\dots O_{nv} | H_i \cdot B \cdot C^n]$. For example, all physicians who employ the same Bayesian network share the same likelihoods, which are built into the network. So they share an empirical context even if each is permitted to input different prior plausibility assessments for disease hypotheses.

Bayes's theorem says that support-function *likelihoods* together with an agent's *prior plausibility assessments* for hypotheses, $P_\alpha[H_i | B]$, determine how strongly for her a hypothesis is supported on the evidence, $P_\alpha[H_i | B \cdot C^n \cdot O_{1u}\dots O_{nv}]$. We might well call this her *posterior plausibility assessment* for the hypothesis on the evidence. I think we have a good handle on where the numerical values of likelihoods or likelihood ratios that contribute to posterior plausibilities are supposed to come from. But perhaps it will be

helpful to say a bit more about how an agent may assess numerical values for *prior plausibilities*.

To understand the numerical values for *priors* more fully, consider the following method by which an agent might make such assessments. She first chooses a hypothesis to use as a basis of comparison. Any hypothesis in the empirical context will do; but a hypothesis she takes to be highly plausible, perhaps the most plausible, will do very nicely. Suppose she chooses H_1 . Then, discounting all information she has that can be *screened off* by possible evidence claims from the evidence partitions in the empirical context,²⁰ but taking into account the weight of all of the various plausibility considerations she regards as relevant, she assesses how many times more or less plausible than H_1 she considers each alternative hypothesis to be (given B). This yields, for each alternative hypothesis H_j , a number r_j for the ratio $P_a[H_j | B]/P_a[H_1 | B]$. And that's it. The prior probabilities themselves are just a way of re-describing these prior plausibility ratio values.²¹ Similarly, posterior plausibilities for hypotheses are just re-descriptions of posterior plausibility ratios. Indeed, the whole Bayesian evaluative enterprise may be understood in terms of probability ratios (as equation 2-4 of section 2 make clear). The posterior plausibility ratios are gotten by multiplying prior plausibility ratios by likelihood ratios.

Agents have a great deal of discretion in assigning prior plausibility (ratio) values. But there are constraints of a sort. If, as evidence accumulates, the likelihood of the evidence on H_i is much greater than on H_j but the agent's posterior plausibility for H_i on this evidence is much smaller than for H_j , then the agent must consider H_i far more plausible prior to the evidence than H_j . Such an assessment may be perfectly legitimate. But, when agents in a community that shares an empirical context disagree enough about priors to resist the evidential influence of the likelihood ratios, they may be expected to give some rationale, some plausibility arguments for their relative plausibility assessments.

Individual agents may not be confident in assessing precise values for priors. But this presents no problem. An individual's imprecise values may be represented by a set of probability functions that covers the range of priors she would find acceptable. The set of priors that represent the variety of plausibility assessments for all members of a community of agents will be larger still. Thus, a community of Bayesian agents may be characterized by a wide-ranging set of prior plausibility assessments for the hypotheses of a context.

Plausibility arguments and the resulting assessments of priors are particularly important when the evidence is not yet sufficiently strong to distinguish among hypotheses, especially when hypotheses appear to be evidentially indistinguishable.²² However, when the true hypothesis differs from its rivals in what it says about the likelihoods of possible observations, an extremely wide range of prior plausibility assessments will very probably be overwhelmed by the accumulating evidence; and the community of agents will be brought to near agreement on posterior plausibility values — near 1 for the true hypothesis and near 0 for its competitors. This will happen even if agents revise their prior plausibilities from time to time — i.e. even if an agent switches from one support function to another due to a reassessment of the prior plausibilities of hypotheses — provided that the true hypothesis is not continually discounted by revising its prior ever downward, approaching zero. The *Likelihood-Ratio Convergence Theorem* discussed in section 2 assures us of this. So, where evidentially distinguishable hypotheses are concerned, assessments of priors may only exert a significant influence in the short run. This is one of the important benefits that accrue to Bayesian agents who employ support functions to govern the updating of their belief-strengths.

6.2 Precisely how support should inform belief

Suppose that the members of a community of Bayesian agents share a common empirical context.

²⁰ In assessing H_j 's *prior* plausibility, any information that can be covered by the evidence partitions must be discounted. Otherwise this information will be double-counted when evidence is brought to bear — both as contributing to the priors and as evidence for hypotheses.

²¹ I.e., the values of the priors are supposed to sum to 1; and for each H_j other than H_1 , $P_a[H_j | B] = r_j \cdot P_a[H_1 | B]$; so we have $1 = \sum_i P_a[H_i | B] = P_a[H_1 | B] \cdot (1 + \sum_j r_j)$; thus, $P_a[H_1 | B] = 1/(1 + \sum_j r_j)$ and $P_a[H_j | B] = r_j/(1 + \sum_j r_j)$.

²² Consider, for example, arguments offered by proponents of various interpretations of quantum mechanics. Such plausibility arguments are seldom conclusive. But they may eliminate some views and provide some basis for agents' plausibility rankings among those that remain.

Intuitively, each agent's *belief strength* for a hypothesis should just be the degree to which the hypothesis is *supported*, according to her *support function*, by the totality of the evidence she has acquired together with whatever prior plausibility considerations she finds relevant. This is the scientific way for an agent to reason — it's the epistemic point of the scientific enterprise.

This idea, that a scientific agent's belief function should be governed in some way by a separate support function, is not entirely new. Carnap (1971) embraces a version of it, and Skyrms (1986) largely concurs, but suggests that the Carnapian account be extended to uncertain evidence. Carnap and Skyrms both describe two-function versions of the now standard Bayesian update model. On this model the ideal agent begins (at birth?) with a support function. The agent's initial belief function is his support function conditional on a tautology: $Q_{\alpha\text{-initial}}[S] = P_{\alpha}[S | A \vee \neg A]$, for all statements S . Carnap seems to think that P_{α} should be a Carnapian logical probability function, but Skyrms apparently sees no need for that. In any case, beginning with $Q_{\alpha\text{-initial}}$, as the agent learns each new bit of evidence E_1, \dots, E_n , he updates his belief function on that evidence as follows: if at a time t this is his total evidence, then his belief function should be $Q_{\alpha,t}[S] = P_{\alpha}[S | E_1 \dots E_n]$; and if at time $t+1$ he becomes certain of new evidence E_{n+1} , his new belief strengths becomes $Q_{\alpha,t+1}[S] = Q_{\alpha,t}[S | E_{n+1}] = P_{\alpha}[S | E_1 \dots E_n \cdot E_{n+1}]$. This is an adaptation to a two-function model of a now fairly standard diachronic Bayesian model of belief updating. Skyrms adds to this the suggestion that if the E_i are not believed with certainty, but the agent possesses some information about each of E_1, \dots, E_n, E_{n+1} , in that order, and possesses no information about other evidence claims, then Jeffrey updating should apply: $Q_{\alpha,t+1}[S] = \sum_{\{E_{n+1}\}} Q_{\alpha,t}[S | E_{n+1}] \cdot Q_{\alpha,t+1}[E_{n+1}]$, where the sum is over E_{k+1} and each of *its* possible alternatives. It turns out that the right way to pound this down to the *founding support function* is this: $Q_{\alpha,t+1}[S] = \sum_{\{E_{n+1}\}} \dots \sum_{\{E_1\}} P_{\alpha}[S \cdot E_1 \dots E_n \cdot E_{n+1}] \cdot (Q_{\alpha-1}[E_1]/P_{\alpha}[E_1]) \cdot \dots \cdot (Q_{\alpha-1}[E_{n+1}]/Q_{\alpha-1}[E_{n+1}])$, (see Hawthorne, 2004a). When S is a hypothesis H for which there are objective or public likelihoods, this updating scheme may also be written as follows: $Q_{\alpha,t+1}[H] = P_{\alpha}[H] \cdot \sum_{\{E_{n+1}\}} \dots \sum_{\{E_1\}} P[E_1 \dots E_n \cdot E_{n+1} | H] \cdot (Q_{\alpha-1}[E_1]/P_{\alpha}[E_1]) \cdot \dots \cdot (Q_{\alpha-1}[E_{n+1}]/Q_{\alpha-1}[E_{n+1}])$.

Pretty as this picture is, it is not the most general account of how support should constrain belief in a model of scientific agents. For one thing, it supposes that each agent must stick with a single *founding support function*, and must only update it in the standard Bayesian or Jeffrey way. This may be an accurate account of how agents usually employ automated Bayesian networks. But it is a poor model of how scientist behave, or should behave. In particular, scientific agents should be free to reconsider their initial plausibility assessments from time to time as new plausibility arguments are introduced, and previous arguments are reevaluated. This seems a natural part of the conceptual development of a science. The Bayesian model of belief revision should accommodate this kind of updating, rather than rule it out of bounds.

Recently Marc Lange (1999) has suggested an alternative way of looking at the standard model of Bayesian updating. Lange's idea is that the standard Bayesian updating scheme is best thought of not as a diachronic account of how agents should update their belief functions over time, but rather as a synchronic model of how an agent should justify his present belief strengths. The idea is that at any given time an agent should be prepared to provide justifying arguments for her present belief strengths for hypotheses. Such arguments should begin by citing a support function, which presumably has public likelihoods and some plausibility arguments for the selection of prior probabilities (though Lange doesn't say much about this). From there a justificatory argument proceeds in a way that looks formally just like standard Bayesian updating. On Lange's account each 'update step', $Q_{\alpha,t+1}[S] = Q_{\alpha,t}[S | E_{n+1}]$, is taken to be a step in an argument that begins with the initial support function and adds in evidence piece by piece until all known relevant evidence has been added. The end result should be the agent's current belief function, which has now been provided a justifying Bayesian argument.

I agree with what I take to be the core idea in Lange's treatment — that a Bayesian support function should be the basis on which an agent's belief strengths for scientific hypotheses are (synchronically) justified. The primary role of a support function in the two-function model I am proposing is to underwrite and justify the agent's belief strengths for hypotheses. However, Lange also takes the step-by-step process of 'updating' by adding in individual bits of evidence to be central to his account of the justificatory argument for an agent's belief function. I see no need to draw on this in my account of the support-belief connection. Rather, as I see it, the justificatory argument may just as well come in a single step. Roughly, if $Q_{\alpha}[C^n \cdot E^n] = 1$ and if this is the total evidence relevant to H_i and its alternatives on which α is

certain, then the agent's belief function should be *aligned* with his support function on this evidence as follows: for each hypothesis, $Q_\alpha[H_j | B \cdot C^n \cdot E^n]$ should equal $P_\alpha[H_j | B \cdot C^n \cdot E^n]$ on this evidence. It then immediately follows that $Q_\alpha[H_j | B] = P_\alpha[H_j | B \cdot C^n \cdot E^n]$; and furthermore, if $Q_\alpha[B] \approx 1$, then $Q_\alpha[H_j] \approx Q_\alpha[H_j | B]$.²³

This idea works fine if all evidence is believed with certainty. But, as I argued earlier, we often possess some evidence for our evidence — i.e. we often possess *uncertain evidence* and *disjunctively-entangled evidence*. So a full account of how support informs and justifies belief should take uncertain evidence into account. Spelling out the details of how this should work is a bit more complicated.

I think it will prove helpful if we first spell out the above idea, which applies only to *certain evidence*, more formally. This should help elucidate the somewhat parallel treatment of *uncertain evidence* that follows. To first approximation a scientific agent's support function should inform her *belief function* according to the following rule:

Approximate Belief-Support Alignment Condition: α is *belief-support aligned* (at a time) for a sequence of evidential conditions C^m of an empirical context *just in case* (at that time):

- (1) α considers one of C^m 's outcome sequences, $(O_{1u} \dots O_{mv})$, to be certain, given $B \cdot C^m$ (i.e. $Q_\alpha[O_{1u} \dots O_{mv} | B \cdot C^m] = 1$), and α 's belief function Q_α agrees with her support function P_α on the posterior probabilities of hypotheses in the empirical context — i.e. $Q_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}] = P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}]$; and
- (2) there is no more inclusive sequence of possible evidence, C^n , such that α considers one of its outcome sequences $(O_{1u} \dots O_{nv})$ to be certain, given $B \cdot C^n$.

This is just a more formal account of the idea stated two paragraphs back. Notice that when this condition is satisfied, the following formula follows:

Approximate Belief-Support Connection Formula: For the C^m on which α is *belief-support aligned* and for its outcome sequence on which α is certain, given $B \cdot C^m$ (i.e. $Q_\alpha[O_{1u} \dots O_{mv} | B \cdot C^m] = 1$), $Q_\alpha[H_i | B \cdot C^m] = P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}]$.

Furthermore, if α is certain of C^m , then $Q_\alpha[H_i | B] = P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}]$; and if α is certain of both B and C^m (or nearly certain of them), then $Q_\alpha[H_i] = P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}]$ (or approximately so).

That would be all there is to it if agents never had partial evidence for evidence statements. But agents seldom find themselves in such an epistemically pristine state. To handle uncertain evidence, the right way for *support* to inform *belief* is roughly this: the agent's belief strength for a hypothesis should be the weighted sum of the degrees to which each *possible evidence sequence* supports the hypothesis, weighted by the agent's belief strengths for each of those possible evidence sequences. I'll now spell out this idea more precisely.

Let us say that an agent α is *belief-support aligned* (at a time) for a given empirical context just in case each possible sufficiently long sequence of outcomes would, if *conditionalized on*, suffice to bring the agent's belief function into agreement with her support function on the posterior probabilities of hypotheses in the empirical context. More formally:

Belief-Support Alignment Condition: α is *belief-support aligned* (at a time) for a sequence of evidential conditions C^m of an empirical context *just in case* (at that time):

- (1) for each of C^m 's outcome sequences $(O_{1u} \dots O_{mv})$ that α considers possible, given $B \cdot C^m$ (i.e. whenever $Q_\alpha[O_{1u} \dots O_{mv} | B \cdot C^m] > 0$), α 's belief function Q_α agrees with her support function P_α on the posterior probability of hypotheses in the empirical context — i.e. $Q_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}] = P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}]$; and
- (2) for each more inclusive sequence of evidence conditions C^n and each of its outcome sequences $(O_{1u} \dots O_{nv})$ that α considers possible, given $B \cdot C^n$, α 's belief function Q_α continues to agree with

²³ This corresponds to what Skyrms (1986) calls the *Certainty Model*.

- her support function P_α on the posterior probabilities of hypotheses in the empirical context — i.e. $Q_\alpha[H_i | B \cdot C^n \cdot O_{1u} \dots O_{nv}] = P_\alpha[H_i | B \cdot C^n \cdot O_{1u} \dots O_{nv}]$.
- (3) no less inclusive subsequence of the evidence conditions in C^m satisfies both conditions (1) and (2).

This *alignment condition* does not presuppose that the relevant possible outcome sequences, $(O_{1u} \dots O_{mv})$, are *believed with certainty* by the agent (though some or all of the outcomes O_{jk} may be certain for α). Rather, the idea is that each of these possible outcome sequences would be informationally complete enough to effectively override all of the evidential information, including partial information, the agent possesses that is relevant to the hypotheses under consideration.

What does it mean for a possible outcome sequence to be *informationally complete enough to override* the agent's evidential information? (i) Suppose the agent is certain of outcome O_{jk} . Then each of the aligning possible outcome sequences considered possible by α must contain O_{jk} . (ii) Suppose the agent is not certain of O_{jk} , but is certain of a disjunctive sentence of form $(O_{jk} \vee D)$, where O_{jk} is considered possible by α . Then among the possible aligning outcome sequences will be those containing O_{jk} ; and they *override* the information ' $(O_{jk} \vee D)$ ' because they entail it. In addition, any sequence containing an alternative to O_{jk} , say O_{jl} , will also *override* ' $(O_{jk} \vee D)$ ' because it entails ' $\neg O_{jk}$ ' (which together with ' $(O_{jk} \vee D)$ ' entails D , which entails $(O_{jk} \vee D)$). (iii) Finally, suppose that the agent has partial information of some other kind for O_{jk} — e.g., evidence for O_{jk} that makes O_{jk} a kind of uncertain evidence for the hypotheses. Then each outcome sequence that contains O_{jk} itself or contains one of its alternatives will override whatever partial information may be relevant to O_{jk} .

Another way of thinking about the general *Belief-Support Alignment Condition* is this. The condition effectively says that *were* the agent to become certain of any one of the outcome sequences of aligning condition C^m , but acquired no other information of relevance to the hypotheses, then her updated belief function, $Q_{\alpha\text{-new}}[*] = Q_\alpha[* | O_{1u} \dots O_{mv}]$, would satisfy the *Approximate Belief-Support Alignment Condition*. The *Approximate Belief-Support Connection Formula* would then follow for that sequence of outcomes.

How should support inform belief under *alignment* when the agent is not certain of any single outcome sequence of the aligning condition C^m ? The obvious idea is that belief should conform to the *expected value* of support — i.e. belief should correspond to the weighted average of the amount of support a hypothesis receives from each possible aligning outcome sequence, weighted by how strongly the agent believes each of the aligning outcome sequences. To see how this works, first notice that it is a theorem of probability theory that $Q_\alpha[H_i | B \cdot C^n] = \sum_u \dots \sum_v Q_\alpha[H_i | B \cdot C^n \cdot O_{1u} \dots O_{nv}] \cdot Q_\alpha[O_{1u} \dots O_{nv} | B \cdot C^n]$. So, when an agent is belief-support aligned with respect to the possible outcomes of C^m , the next formula follows as a theorem:

The Belief-Support Connection Formula: For the C^m on which α is *belief-support aligned*,
 $Q_\alpha[H_i | B \cdot C^m] = \sum_u \dots \sum_v P_\alpha[H_i | B \cdot C^m \cdot O_{1u} \dots O_{mv}] \cdot Q_\alpha[O_{1u} \dots O_{mv} | B \cdot C^m]$.
 Furthermore, if α is certain of C^m , then $Q_\alpha[H_i | B] = Q_\alpha[H_i | B \cdot C^m]$; and if α is certain of both B and C^m (or nearly certain of them), then $Q_\alpha[H_i] = Q_\alpha[H_i | B \cdot C^m]$ (or approximately so).²⁴

When *alignment* holds, the *belief-support connection formula* is a generalization of the traditional Bayesian *total evidence requirement* to cases where at least some of the evidence may be uncertain evidence. More specifically, the usual total evidence requirement says that an agent's belief strength for a hypothesis should match the evidential support for the hypothesis by the totality of the agent's *certain* evidence. And that is just what the *approximate belief-support connection formula* says. This would be right if, for each possible condition C_k , the agent is either certain of one of its possible outcomes, or has no information at all relevant to its outcomes. But agents often possess some inconclusive bits of evidence (or even just hunches) that alter their belief strengths in some of the possible outcomes of experiments or observations. When they do, the total evidence requirement as usually expressed misses out on some of the evidence that influences the agent's belief strengths. *Alignment* together with the general *connection formula* extends the usual idea. It tells us that an

²⁴ More generally, if $B \cdot C^n$ is uncertain, then there may be partitions of alternatives $\{B_v\}$ and $\{C_u^n\}$. In that case we may have, $Q_\alpha[H_i] = \sum_v \sum_u Q_\alpha[H_i | B_v \cdot C_u^n] \cdot Q_\alpha[B_v \cdot C_u^n]$.

agent's belief strength for a hypothesis should match the evidential support for the hypothesis by the *totality* of the agent's evidence, regardless of how certain that evidence is.

6.3 The old evidence problems dissolve

The recognition of support functions as distinct from current belief functions, and of the connection of the two through alignment, dissolves the problems with old evidence. For, even if the agent is certain of E , so her belief function has $Q_\alpha[E] = 1$, it is her support function that measures the degree to which E supports H_i (given $B \cdot C$). And her support function will not assign E support-strength 1 unless H_i and all of its alternatives yield degree-1 likelihoods for E . Thus, support functions provide a measure that satisfies the Bayesian intuitions about the basis of evidential support in public or objective likelihoods, and overcomes the old evidence problems that plague a Bayesianism that restricts itself entirely to current belief-strengths.

To see more precisely how the old evidence problems dissolve, consider the way in which Bayesians assess the degree to which individual pieces of evidence favor a hypothesis. To this end Bayesians have developed several measures of the incremental confirmation. Among the most prominent are the difference measure, $Q_\alpha[H | B \cdot C \cdot E] - Q_\alpha[H | B]$, the ratio measure, $Q_\alpha[H | B \cdot C \cdot E] / Q_\alpha[H | B]$, and the likelihood ratio measure, $Q_\alpha[E | H \cdot B \cdot C] / Q_\alpha[E | \neg H \cdot B \cdot C]$. However, whenever the belief strength of E changes, bringing about an update of Q_α to a new belief function $Q_{\alpha\text{-new}}$, each of these measures of incremental confirmation, re-stated in terms of $Q_{\alpha\text{-new}}$, takes on a new value as well. For, each of these measures is sensitive to the agent's current belief strength for E . Presumably, though, the amount of incremental confirmation that E affords H_i (given $B \cdot C$) should not be sensitive to how strongly the agent happens to believe E at the moment. Thus, couched in terms of current belief functions, none of these measures of incremental confirmation works properly.

This problem is easily overcome if we recognize the existence of support functions. Intuitively it makes better sense for incremental confirmation to be measured in terms of how much a specific evidential claim supports a hypothesis, not what the agent's belief strengths regarding hypotheses and evidence claims happen to be at a given time. And agents' support functions remain constant through evidential updating. So, when incremental confirmation is measured in terms of an agent's support function, each of the proposed measures of incremental confirmation suggested above (with ' P_α ' in place of ' Q_α ') remains constant across belief updates.²⁵

Taking support functions to be distinct from current belief functions also explains how the recognition of the appropriate public or objective degree-of-support likelihoods can result in new support for the hypothesis by evidence already known. Before these likelihoods are recognized, the empirical import of H for E remains vague for the agent. Formally, E does not belong to any of the evidence partitions included in the community's present empirical context, and the various support functions represented in the context may disagree on the likelihood of E . When the objective or public likelihoods of E and its alternatives are acknowledged by the scientific community, the empirical context is expanded to include a new evidence partition, and only those support functions that take on the agreed public values for the likelihoods for that new partition are retained within the new empirical context. Then, alignment with the (restricted) support function for the extended empirical context, now informed by E , may change the agent's belief strength for H .²⁶

²⁵ See (Fitelson, 1999) for an excellent comparative analysis of various incremental confirmation measures.

²⁶ In cases where the likelihoods for already known E are strictly logical (i.e. where either $H_i \cdot B \cdot C$ logically entails E or where some proper account of *direct inference* completely specifies the appropriate likelihoods), the support function in the old empirical context must already specify precise values for the likelihoods. Even so, E and its alternatives may have escaped inclusion in the collection of evidence partitions in the empirical context because the relevance of E was not recognized — the logical relationship was overlooked. In such cases belief-support alignment based on that incomplete empirical context will fail to accord with the agent's total available evidence, since E has been left out. Then, recognition of the logical relationship should bring about expansion to a new empirical context that

6.4 The benefits of alignment

Alignment provides a plausible model of the belief-support connection. And certain benefits accrue to agents who align in a shared empirical context.

Consider a community of agents who share an empirical context. Through the usual scientific process of acquiring and reporting evidence (authenticated through the usual review process), each of these agents may legitimately come to very strongly believe a sequence of true evidence claims that goes well beyond the evidence he might acquire on his own. Furthermore, Bayesian convergence results (such as the *Likelihood Ratio Convergence Theorem* described in section 2) say that as evidence accumulates, it becomes extremely likely that the true hypothesis will beat out each of its empirically distinct competitors in a contest of degree-of-support likelihood ratios. When this happens, the degree of support for false alternatives must fall towards zero, and the degree of support of the disjunction of the true hypothesis with its empirically equivalent rivals approaches 1. As a result, each *aligned agent* will come to strongly believe that the false empirically distinct alternatives are indeed false. And each such agent will come to strongly believe that the disjunction of the true hypothesis with its empirically equivalent rivals is true. Indeed, each *aligned agent* may come to strongly believe the true hypothesis itself if the community can find sufficiently probative plausibility considerations to lay low its empirically equivalent rivals. Thus, to the extent that an agent benefits from strongly believing that false hypotheses are false and that true hypotheses are true, there is good reason for him to align *belief* with *support*.²⁷

JAMES HAWTHORNE

Department of Philosophy
University of Oklahoma
605 Dale Hall Tower
455 West Lindsey Street
Norman, OK 73019
USA
hawthorne@ou.edu

Appendix

Theorem: We restrict our attention to those hypotheses H_k and possible observational outcomes O_i (from their respective partitions) for which α 's belief strengths $Q_\alpha[H_k]$ and $Q_\alpha[O_i]$ are greater than 0.

Suppose α 's belief updates from Q_α to $Q_{\alpha\text{-new}}$ don't outright falsify any of these H_k or O_i — i.e., $Q_{\alpha\text{-new}}[H_k] > 0$ and $Q_{\alpha\text{-new}}[O_i] > 0$ (as is appropriate to updating with uncertain evidence). Suppose

includes a partition containing E . This results in a new belief-support alignment relationship that better accounts for agents' appropriate belief-strengths for the hypotheses.

²⁷ I am indebted to many friends and colleagues for helping me work through the ideas presented here. I especially thank Luc Bovens, Bas Van Fraassen, Adam Morton, Chris Swoyer, and the referees at *Mind*. I presented an early draft at the 31th Annual Meeting of the Society for Exact Philosophy, in Montreal, in May of 2001. Thanks to the participants for their helpful comments. Part of a sabbatical leave granted by the University of Oklahoma provided me the opportunity to do additional work on this project in the winter of 2002-2003, during a three-month stay with the research group on *Philosophy, Probability, and Modeling* at the University of Konstanz (supported by the Alexander von Humboldt Foundation, the Federal Ministry of Education and Research and the Program for Investment in the Future (ZIP) of the German Government through a Sofja Kovalevskaja Award). Thanks to the participants and research fellows in the group who discussed Bayesian inference with me, especially Stephan Hartmann, Franz Dietrich, and Franz Huber.

also that for each such O_i and H_k , $Q_{\alpha\text{-new}}[H_k | O_i] = Q_{\alpha}[H_k | O_i]$ (which is the Jeffrey *rigidity* condition for updating on uncertain evidence). And suppose there is an O_i such that $Q_{\alpha\text{-new}}[O_i] \neq Q_{\alpha}[O_i]$ (i.e., there's a change in the belief strength of at least one of the possible outcomes). Then, provided that there is at least one O_g to which every O_i is *connected* (as specified below, relative to Q_{α} and partition $\{H_k\}$), for some O_i and H_k , the likelihood must change — $Q_{\alpha\text{-new}}[O_i | H_k] \neq Q_{\alpha}[O_i | H_k]$.

Notice, we are not supposing that the likelihoods $Q_{\alpha\text{-new}}[O_i | H_k]$ and $Q_{\alpha}[O_i | H_k]$ are nonzero, nor are we supposing that the posterior probabilities $Q_{\alpha\text{-new}}[H_k | O_i]$ and $Q_{\alpha}[H_k | O_i]$ are non-zero. We assume only that for those H_k and O_i for which $Q_{\alpha}[H_k] > 0$ and $Q_{\alpha}[O_i] > 0$, their updates on uncertain evidence, $Q_{\alpha\text{-new}}[H_k]$ and $Q_{\alpha\text{-new}}[O_i]$, remain greater than 0. Throughout the following discussion and proof of the theorem we restrict our attention to those hypotheses H_k and possible outcomes O_i for which α 's belief strengths $Q_{\alpha}[H_k]$ and $Q_{\alpha}[O_i]$ are greater than 0.

The notion I call *connectedness* between pairs of outcomes O_i and O_j (relative to the hypotheses H_k for function Q_{α}) is defined in terms of the notion of *degree of separation* between O_i and O_j , which is itself defined as follows:

- (read ' $[O_i \text{ sep}_n O_j]$ ' as 'there are no more than n degrees of separation between O_i and O_j)
- (0) $[O_i \text{ sep}_0 O_j]$: O_j just is O_i , and for some H_k , $Q_{\alpha}[O_i | H_k] > 0$ (this is a special case);
 - (1) $[O_i \text{ sep}_1 O_j]$: for some H_k , $Q_{\alpha}[O_i | H_k] > 0$ and $Q_{\alpha}[O_j | H_k] > 0$;
 - ($n > 1$) for each $n > 1$, $[O_i \text{ sep}_n O_j]$: for some O_g , $[O_i \text{ sep}_{n-1} O_g]$ and $[O_g \text{ sep}_{n-1} O_j]$;
 - (∞) O_i is *connected* to O_j (i.e. O_i is not completely separated from O_j) : for some n , $[O_i \text{ sep}_n O_j]$.

Notice that since we are supposing $Q_{\alpha}[O_i] > 0$, for each O_i under discussion, it will always be the case that $Q_{\alpha}[O_i | H_k] > 0$ for some H_k under discussion (i.e. such that $Q_{\alpha}[H_k] > 0$): because $Q_{\alpha}[O_i] = \sum_k Q_{\alpha}[O_i | H_k] \cdot Q_{\alpha}[H_k]$. So Clause (0) is automatically satisfied — all of the O_i are sep_0 from themselves. However, satisfaction of Clause (1), of $[O_i \text{ sep}_1 O_j]$, is not automatic, since it is generally possible that on each H_k for which $Q_{\alpha}[O_i | H_k] > 0$, $Q_{\alpha}[O_j | H_k] = 0$. For that to happen we would have to have the following situation: if either of the two possible outcomes occurred, and α became certain of it, all hypotheses compatible with the other possible outcome would become falsified. But even if $[O_i \text{ sep}_1 O_j]$ were to fail, there may well be two hypotheses, H_u and H_v , such that $Q_{\alpha}[O_i | H_u] > 0$ and $Q_{\alpha}[O_j | H_v] > 0$, and such that each gives a positive likelihood to some third possible outcome, O_g . In that case we would at least have $[O_i \text{ sep}_2 O_j]$.

It is easy to establish (by induction on n) that the sep_n are cumulative (i.e., if $[O_i \text{ sep}_{n-1} O_j]$, then, because $[O_j \text{ sep}_{n-1} O_j]$, we have $[O_i \text{ sep}_n O_j]$). It is also easy to see (by induction) that each sup_n is reflexive and symmetric. So *connectedness* is also reflexive and symmetric. And although sup_n need not be transitive, *connectedness* is transitive: i.e., suppose O_i is *connected* to O_j and O_j is *connected* to O_k ; then for some n , $[O_i \text{ sep}_n O_j]$ and $[O_j \text{ sep}_n O_k]$; but then $[O_i \text{ sep}_{n+1} O_k]$. Thus, *connectedness* is an equivalence relation.

Intuitively, *connectedness* works like this. Imagine the hypotheses H_k lined up horizontally, and imagine the possible outcomes O_i lined up some distance beneath them. Draw a link from each hypothesis to each of the outcomes to which it assigns a non-zero likelihood. Now, from any given outcome O_i try to draw a continuous path to another outcome O_j by following the links (up to a hypothesis, down to another outcome, up to another hypothesis, down to another outcome, etc.). If one can draw a path from O_i to O_j in this way, then O_i is *connected* to O_j . For example, when $[O_i \text{ sep}_1 O_j]$ holds, only two links are required, one from O_i to some H_k and one from that H_k to O_j . Notice that if some H_k assigns non-zero likelihood to every O_i , then each pair of outcomes is sep_1 related; so all O_i are *connected*. Similarly, if some O_i is assigned non-zero likelihood by every H_k , then each pair of outcomes is sep_1 related; so all O_i are *connected*. In general, sup_n related outcomes might only be linked through some more complicated route. If the shortest route between O_i and O_j employs $n+1$ links, then $[O_i \text{ sep}_{n-1} O_j]$ fails, but $[O_i \text{ sep}_n O_j]$ succeeds.

Proof of the Theorem: From this point forward let us assume that all outcomes are *connected* — i.e. there is an O_j such that for each O_i , O_i is *connected to* O_j . This condition will clearly be satisfied in many confirmational contexts in the sciences. Let us also suppose that updating maintains the same support for hypotheses on each given O_i (i.e. $Q_{\alpha\text{-new}}[H_k | O_i] = Q_{\alpha}[H_k | O_i]$) and maintains the same values for likelihoods (i.e. $Q_{\alpha\text{-new}}[O_i | H_k] = Q_{\alpha}[O_i | H_k]$). We will see that under these conditions, $Q_{\alpha\text{-new}}[O_i] = Q_{\alpha}[O_i]$ for all O_i — so no updating of the O_i can have occurred after all.

The proof proceeds by first showing that if O_i is *connected to* O_j , then $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[O_j]/Q_{\alpha}[O_j]$. This follows from the fact that for each n , if $[O_i \text{ sep}_n O_j]$, then $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[O_j]/Q_{\alpha}[O_j]$, which is proved by induction, as follows:

It clearly holds for sep_0 , but sep_0 is a special case. So let's have our basis step be for sep_1 . Suppose $[O_i \text{ sep}_1 O_j]$. Then for some H_k , $Q_{\alpha}[O_i | H_k] > 0$ and $Q_{\alpha}[O_j | H_k] > 0$. But notice that, for every H_m and all O_g , $Q_{\alpha\text{-new}}[O_g | H_m] \cdot Q_{\alpha\text{-new}}[H_m]/Q_{\alpha\text{-new}}[O_g] = Q_{\alpha\text{-new}}[H_m | O_g] = Q_{\alpha}[H_m | O_g] = Q_{\alpha}[O_g | H_m] \cdot Q_{\alpha}[H_m]/Q_{\alpha}[O_g]$. Then, since $Q_{\alpha\text{-new}}[O_i | H_k] = Q_{\alpha}[O_i | H_k] > 0$, we have $Q_{\alpha\text{-new}}[H_k]/Q_{\alpha\text{-new}}[O_i] = Q_{\alpha}[H_k]/Q_{\alpha}[O_i]$ — i.e. we have $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[H_k]/Q_{\alpha}[H_k]$. Similarly, since $Q_{\alpha\text{-new}}[O_j | H_k] = Q_{\alpha}[O_j | H_k] > 0$, we have $Q_{\alpha\text{-new}}[O_j]/Q_{\alpha}[O_j] = Q_{\alpha\text{-new}}[H_k]/Q_{\alpha}[H_k]$. Thus, $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[O_j]/Q_{\alpha}[O_j]$.

The induction step is quite easy. Suppose the relationship holds for sup_n and suppose $[O_i \text{ sep}_{n+1} O_j]$. Then for some O_g , $O_i \text{ sep}_n O_g$ and $O_g \text{ sep}_n O_j$; so $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[O_g]/Q_{\alpha}[O_g] = Q_{\alpha\text{-new}}[O_j]/Q_{\alpha}[O_j]$.

Now, given our assumption that there is some O_g to which all O_i are connected, we have, for all O_i , $Q_{\alpha\text{-new}}[O_i]/Q_{\alpha}[O_i] = Q_{\alpha\text{-new}}[O_g]/Q_{\alpha}[O_g]$. So, $Q_{\alpha\text{-new}}[O_i] \cdot Q_{\alpha}[O_g] = Q_{\alpha\text{-new}}[O_g] \cdot Q_{\alpha}[O_i]$, for all O_i . Summing both sides of this equation over all O_i yields $Q_{\alpha}[O_g] = Q_{\alpha\text{-new}}[O_g]$. Factoring this equivalence out of the previous equation then yields $Q_{\alpha\text{-new}}[O_i] = Q_{\alpha}[O_i]$, for every O_i .

References

- Bogdan, Radu J. (ed) 1981: *Henry E. Kyburg, Jr. & Isaac Levi*. Dordrecht: D. Reidel.
- Carnap, Rudolf, 1971: 'Inductive Logic and Rational Decisions', in Carnap and Jeffrey 1971, pp. 5–31.
- Carnap, Rudolf and Richard Jeffrey (eds) 1971: *Studies in Inductive Logic and Probability I*. Los Angeles: University of California Press.
- Earman, John 1992: *Bayes or Bust?* Cambridge, MA: MIT Press.
- Edwards, W., H. Lindman, L. J. Savage 1963: 'Bayesian Statistical Inference for Psychological Research'. *Psychological Review*, 70, pp. 193–242.
- Eells, Ellery 1985: 'Problems of Old Evidence'. *Pacific Philosophical Quarterly*, 66, pp. 283–302.
- Field, Hartry 1977: 'Logic, Meaning, and Conceptual Role'. *Journal of Philosophy*, 74, pp. 379–409.
- Fitelson, Branden 1999: 'The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity'. *Philosophy of Science*, 66 (Proceedings), pp. 362–78.
- Glymour, Clark 1980: *Theory and Evidence*. Princeton: Princeton University Press.
- Hajek, Alan 2003: 'What Conditional Probability Could Not Be'. *Synthese*, 137, pp. 273–323.
- Harper, William L. 1981: 'Kyburg on Direct Inference', in Bogdan 1981, pp. 97–127.
- Hawthorne, James 1993: 'Bayesian Induction Is Eliminative Induction'. *Philosophical Topics*, 21, pp. 99–138.
- 2004a: 'Three Models of Sequential Belief Updating on Uncertain Evidence'. *Journal of Philosophical Logic*, v. 33, pp. 89–123.
- 2004b: 'Inductive Logic'. *The Stanford Encyclopedia of Philosophy (Fall 2004 Edition)*, Edward N. Zalta (ed), URL = <<http://plato.stanford.edu/archives/fall2004/entries/logic-inductive/>>.
- Howson, Colin and Peter Urbach 1993: *Scientific Reasoning: the Bayesian Approach*, 2nd ed. Chicago: Open Court.
- Jeffrey, Richard 1965: *The Logic of Decision*. New York: MacGraw-Hill.
- 1970: 'Dracula Meets Wolf Man: Acceptance vs. Partial Belief', in Swain 1970, pp. 157–85.

- 1988: 'Conditioning, Kinematics, and Exchangeability', in Skyrms and Harper 1988. Reprinted in Jeffrey 1992.
- 1991: 'Bayesianism with a Human Face', in Jeffrey 1992.
- 1992: *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.
- Kyburg, Henry 1977: 'Randomness and the Right Reference Class'. *Journal of Philosophy*, 74, pp. 501–20.
- Lange, Marc 1999: 'Calibration and the Epistemological Role of Bayesian Conditionalization'. *Journal of Philosophy*, 96, pp. 294-324.
- Levi, Isaac 1977: 'Direct Inference'. *Journal of Philosophy*, 74, pp. 5–29.
- 1978: 'Confirmational Conditionalization'. *Journal of Philosophy*, 75, pp. 730–7.
- 1980: *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
- Popper, Karl R. 1959: *The Logic of Scientific Discovery*. New York: Basic Books.
- Skyrms, Brian 1986: *Choice and Chance*, 3rd edition. Belmont, CA: Wadsworth.
- Skyrms, Brian and William Harper (eds) 1988: *Causation, Chance, and Credence*, Vol. 1. Dordrecht: Kluwer.
- Swain, M. (ed) 1970: *Induction, Acceptance and Rational Belief*. Dordrecht: D. Reidel.