Philosophy of Science

June, 1996

SUBJECTIVE AND OBJECTIVE CONFIRMATION*

PATRICK MAHER†

Department of Philosophy
University of Illinois at Urbana-Champaign

Confirmation is commonly identified with positive relevance, E being said to confirm H if and only if E increases the probability of H. Today, analyses of this general kind are usually Bayesian ones that take the relevant probabilities to be subjective. I argue that these subjective Bayesian analyses are irremediably flawed. In their place I propose a relevance analysis that makes confirmation objective and which, I show, avoids the flaws of the subjective analyses. What I am proposing is in some ways a return to Carnap's conception of confirmation, though there are also important differences between my analysis and his. My analysis includes new accounts of what evidence is and of the indexicality of confirmation claims. Finally, I defend my analysis against Achinstein's criticisms of the relevance concept of confirmation.

By *confirmation* I mean the relation that holds between E and H when E is evidence for H. There is considerable prima facie plausibility to the view that this relation holds iff E makes H more probable, i.e., iff $p(H|E) > p(H)$ for a suitable probability function p. I will call this view *the relevance concept of confirmation.* As stated, this concept is vague, especially as the "suitable probability function" has not been specified. In this paper I will discuss different attempts at giving a precise formulation of the relevance concept of confirmation; I will call such formulations *relevance analyses* of confirmation.

For a couple of decades, Carnap's *Logical Foundations of Probability* (1950) provided the standard relevance analysis of confirmation; it was an analysis that made confirmation an objective relation (Section 1). However, difficulties with Carnap's program in confirmation theory led to the

*Received March 1995.

†Send reprint requests to the author, 105 Gregory Hall, 810 S. Wright Street, Urbana, IL 61801.

Philosophy of Science, 63 (June 1996) pp. 149–174. 0031-8248/96/6302-0001$2.00

subsequent development of Bayesian relevance analyses that view confirmation as subjective. In this paper I argue that these subjective analyses are all untenable (Section 2). I offer in their place a relevance analysis that is like Carnap's in making confirmation objective but which avoids what I agree to be the untenable features of his analysis, as well as avoiding the flaws of the subjective analyses (Section 3). Finally, I defend this analysis against Achinstein's criticisms of the relevance concept of confirmation (Section 4).

1. Carnap. Carnap distinguished what he called classificatory, comparative, and quantitative concepts of confirmation; it is the first of these that is my concern in this essay. Carnap says that this classificatory concept of confirmation is expressed by such sentences as '*E* gives some (positive) evidence for *H*' (1950, 21), which suggests that his classificatory concept of confirmation is the same as confirmation in my sense.

However, in one respect Carnap's conception of confirmation deliberately departs from the ordinary notion of what it is for *E* to be evidence for *H*. This departure is that Carnap allows *E* to confirm *H* regardless of whether *E* is evidence or even known to be true. This was a deliberate decision, motivated by Carnap's view that his task was to identify a relation, analogous to entailment, such that no factual knowledge is needed to know whether or not *E* confirms *H* (1950, 20f.). Carnap (1950, 468) acknowledged that he was following Hempel (1945, 24) on this point. Both authors are aware that in this respect their usage departs from ordinary language.

Carnap's analysis of confirmation assumes that there is a unique probability function c such that $c(H|E)$ represents the probability that it is rational to give to *H* when the total evidence is *E*.[1] He gives two analyses of confirmation (1950, 463). One is for a relative notion of confirmation: *E* is said to confirm *H* relative to $\mathcal{B}$ iff $c(H|\mathcal{B}E) > c(H|\mathcal{B})$. The other is for an absolute concept: *E* is said to confirm *H* "so to speak absolutely; that is, without reference to any prior factual evidence" iff $c(H|E) > c(H|T)$, where *T* is a tautology.

This analysis can be correct only if there really is a unique probability function c with the properties Carnap ascribes to it. However, arguments to establish the existence of this function have not been convincing. The usual attempts to demonstrate its existence have been a variation on the traditional idea that, in the absence of any evidence, a rational person would give all possibilities the same probability. A problem with this is that possibilities can be individuated in different ways, in which case this principle leads to different probability functions, and no convincing

[1]So that the notation in this essay can be uniform I am here departing in several respects from the notation that Carnap used in (1950).

method of identifying the "right" way of individuating the possibilities has yet been proposed. Carnap himself eventually withdrew the claim that there was a uniquely correct probability function c (1971a, 27). He did not, however, provide an analysis of confirmation that did not depend on this retracted claim.

Skepticism about the existence of c led to the development of subjective Bayesian analyses of confirmation; these have dominated post-Carnapian confirmation theory because they do not require there to be a function such as Carnap's c.

2. Subjective Bayesian Analyses. Bayesians hold that a rational person has degrees of certainty that are representable by a probability function; this requirement is called *coherence.* The probability function that represents a person's degrees of certainty is referred to as that person's probability function and its values are said to be the person's probabilities.[2]

In addition to coherence, Bayesians also endorse principles about how a person's probabilities should be changed over time. The original principle of this kind is *conditionalization,* which says that if a person with probability function p learns E and nothing else then that person's posterior probability function should be the result of conditioning p on E, i.e., it should be $p(\cdot|E)$.

A person can satisfy these conditions yet still have degrees of certainty that seem crazy to most of us. Nevertheless, *subjective* Bayesians do not endorse any additional constraints on rational degrees of certainty beyond those just mentioned. This is not to say that subjective Bayesians deny the existence of any further conditions, but they do confirmation theory without assuming the existence of any such further conditions. The rationale for this stance is partly an attitude of tolerance and partly the view that nobody has yet produced any very convincing additional principles of rationality. Since subjective Bayesianism is the dominant form of Bayesianism today, the term 'Bayesian' now commonly means 'subjective Bayesian'.

Subjective Bayesians have offered analyses of what they call confirmation, according to which confirmation is relative to persons and times. However, we do not ordinarily think that what is evidence for what is subjective in this way. If I say that E is evidence for H and you say it is not, we normally suppose that we are contradicting one another. On the subjective Bayesian analyses, my statement is presumably to be interpreted as meaning that E confirms H for me now and your statement as meaning

[2]It is now usual to allow that a person's degrees of certainty may be indeterminate, in which case the person's degrees of certainty would be represented by a *set* of probability functions. This is a complexity that I will put aside; including it would not change my conclusion and Bayesian analyses of confirmation have themselves ignored this point.

that E does not confirm H for you now; these statements are not contradictory.

Some Bayesians think that the idea of objective confirmation is an illusion; Howson has dismissed it as merely "an ancient habit of thought that dies hard" (Howson 1991, 550). In any case, the subjective Bayesian analyses of confirmation cannot charitably be supposed to be analyses of an objective notion of confirmation. But then, since confirmation ordinarily is understood in this objective way, we have to ask what concept Bayesian analyses are meant to be an analysis of. We cannot evaluate the adequacy of a proposed analysis unless we know the concept to which it is meant to correspond. Unfortunately, the Bayesian confirmation literature gives us no guidance on this question.

A natural suggestion is that subjective Bayesian analyses of confirmation are trying to say what it is for a person X at time t to take E to be evidence for H. People clearly do take things to be evidence for other things and we can identify this psychological state without assuming (or denying) that there is any objective relation of confirmation. However, this suggestion cannot be right, for people do not take E to be evidence for anything unless E is actually a piece of evidence, whereas subjective Bayesian analyses allow that, for X at t, E can be evidence for H even if X at t is confident that E is false. (In this regard, as in many others, Bayesian analyses of confirmation show their Carnapian roots.) It therefore appears to me that subjective Bayesian analyses of confirmation are best understood as trying to say what it is for a person to take E to be such that, *if* it were evidence, it *would* be evidence for H. At any rate, they are trying to analyze a concept which is such that, if a person takes E to be evidence, then the concept in question applies to the person iff the person takes E to be evidence for H. I will now evaluate subjective Bayesian analyses of confirmation by seeing how adequate they are as analyses of this psychological concept.

2.1. Simple Bayesianism. The simplest Bayesian analysis of confirmation says that E confirms H for person X at time t iff $p(H|E) > p(H)$, where p is X's probability function at t.

It is now generally acknowledged that the simple Bayesian analysis of confirmation is inadequate, for a reason that Glymour (1980, 85–93) drew attention to: If I am certain of E at t and p is my probability function at t then $p(E) = 1$ and so $p(H|E) = p(H)$, for all H. Thus the simple Bayesian analysis implies that I cannot take E to be evidence for anything. This conclusion is wrong; there is no inconsistency between being certain of E and taking E to be evidence for H.[3]

[3]There is even a prima facie plausible argument that anything taken to be evidence must be taken to be certain: If E is not certain, then what one has learned from experience is not

Bayesians, as well as their critics, have generally agreed that this objection shows the simple Bayesian analysis to be inadequate. Nevertheless, one does occasionally meet with attempts to defend the simple Bayesian analysis against this objection, by saying that evidence never is completely certain, or should not be taken to be completely certain. However, the objection as I have presented it does not assume that evidence ever is or should be completely certain; the key point is rather that there is no inconsistency between E being certain and E being regarded as evidence for H, and the simple Bayesian analysis implies otherwise.[4]

There are two ways of modifying the simple Bayesian analysis that have been endorsed by subjective Bayesians in order to avoid Glymour's objection to that analysis. I will discuss them in the next two subsections.

2.2. Historical Bayesianism. One popular suggestion has been that, when X knows E at t, the probability function to use for judging confirmation is, not the one X has at t, but rather the one X had immediately prior to learning E. I will call this the *historical analysis.* If we assume that, at the time of learning E, X updated by conditioning on E, then another way to express this analysis is that evidence E confirms H for X iff X's probability for H went up at the time that X acquired the evidence E.

One problem with the historical analysis is that H might not have been formulated when E was learned, in which case H had no probability at that time and, *a fortiori,* the probability of H did not go up when E was learned. I will return to this problem shortly, but I will first consider the situation in which the historical analysis is most at home: The case where H was formulated, and given a probability, before E was learned. Horwich (1982, 52f.), Eells (1985, 287), and Earman (1992, 122) have all supposed that, in this case at least, the historical analysis is satisfactory. However, I will show that it is untenable even in this most favorable case.

It will be helpful to have a concrete example of confirmation before us. An example that suits my purposes as well as any is one that has often been cited in the literature of confirmation theory: The confirmation of Einstein's General Theory of Relativity by the advance of the perihelion of Mercury. By the end of the nineteenth century it was known that the perihelion of Mercury was advancing at a rate of 574 arc-seconds per

adequately represented by E. Thus it has been suggested that, in cases where the content of what is learned from experience is not representable by any proposition of which we are certain, the evidence should be viewed as non-propositional (Jeffrey 1965, Pearl 1990). But against this it may be urged that we often do cite as evidence propositions which, when pressed, we should concede are not absolutely certain. In Section 3.1 I will propose a resolution of this conflict that does not require evidence to be absolutely certain.

[4]Earman (1992, 120f.) argues that even when evidence is uncertain the simple Bayesian analysis does not lend itself to an adequate measure of *strength* of confirmation. In this paper, though, I am discussing only analyses of qualitative confirmation.

century. Calculations using Newtonian gravitational theory implied an advance of only 531 arc-seconds per century. In 1915 Einstein found that his General Theory of Relativity implied the observed value of the advance of the perihelion of Mercury. Both the Newtonian and relativistic calculations assumed certain facts about the other members of the solar system. In particular, they assumed that the sun is roughly spherical. In the 1960s, measurements of the sun's shape suggested that the sun is actually sufficiently oblate to reduce the advance of the perihelion of Mercury by 3 arc-seconds per century, in which case the General Theory would no longer agree with the observed value. Thus some advocates of alternatives to Einstein's theory claimed that the advance of the perihelion of Mercury was really evidence against Einstein's theory, not evidence for it (Will 1986).

Let G be the General Theory of Relativity, A the proposition that the perihelion of Mercury advances by 574 arc-seconds per century, as S the proposition that departures of the sun from perfect sphericity are too small to have any observable effect on the perihelion of Mercury. In order to ensure the straightforward applicability of the historical analysis, imagine that before learning either A or S I had a probability function p that was defined for G, A, and S. Suppose that at this time I also knew the other facts necessary to derive A from G and S. Since S and A do not separately provide any reason to believe G, let us also suppose that $p(G|AS) > p(G|S) = p(G|A) = p(G)$. If I learned first S and then A the historical analysis says that A is (for me now) evidence for G while S is not. On the other hand, if I learned these facts in the reverse order then the historical analysis says that their evidential status is also reversed; now S (for me) confirms G while A does not. There are two things wrong with this. First, I can perfectly well judge that A is evidence for G, even if I learned A before learning S. Second, people's judgments of what is evidence for what are normally independent of the order in which facts were learned, contrary to what the historical analysis implies.[5]

Now let me make the example more like the historical situation. Let us suppose that S initially seemed very plausible to me, so that when I learned A my probability for G went up, but later evidence convinced me that S was probably false. I should then in the end judge A to be evidence against G, not evidence for it (just as some physicists did in the 1960s). However, the historical analysis is incompatible with this; it looks only at the probability function I had just prior to learning A, and if q is that probability function, $q(G|A) > q(A)$, so at all later times I am supposed to judge that A confirms G.

[5]It is easy to generate other examples in which the historical analysis goes similarly astray. The interested reader might like to reflect on how the analysis handles cases in which $p(H|E_1) = p(H|E_2) > p(H) > p(H|E_1E_2)$. (Carnap 1950, 382f., gives an example in which the probabilities satisfy these inequalities.)

These examples show that the historical analysis does not correctly capture the concept it is aimed at even in the most favorable case, where the required historical probabilities are defined. This basic problem has not been addressed in the literature. Instead attention has been focused on trying to deal with the case in which the required historical probabilities are not defined. It seems to have been hoped that this case could be dealt with in the following manner: Suppose that (like many physicists active around 1915, including Einstein himself) I knew A before I had thought of G. After I have come to consider G I find that G (together with other plausible assumptions) entails A, and at that time my probability for G goes up. Thus although my probability for G did not go up when I learned A, it did go up when I learned that G (together with the other assumptions) entails A; a modified historical analysis could say that the latter fact captures the content of my judgment that A confirms G.

This modification does not apply to cases in which a person knows the relevance of a piece of evidence to a theory from the moment of first thinking of the theory; for example, students might know that G successfully explained A before they even know what G is (Earman 1992, 130). A more fundamental problem is that, even if my probability for G did go up when I realized that it explained A, it does not follow that I must forever afterwards judge that A confirms G; if I subsequently find that S is false, and that G together with $\bar{S}$ is inconsistent with A, then I will judge that A disconfirms G, contrary to what the historical analysis implies. Thus the attempt to make the historical analysis generally applicable merely enlarges the class of counterexamples to it.

2.3 Counterfactual Bayesianism. An alternative to the historical analysis has been to invoke counterfactual conditionals. In its simplest form, this approach says that, when X knows E at t, the probability function to use for judging confirmation is the one X *would* have *if* X did not know E. This counterfactual analysis was given a qualified endorsement by Garber (1983, 103). Howson (1984, 1985, 1991) has defended a version of the analysis that is more difficult to understand but which seems to amount to essentially the same thing.[6]

The counterfactual analysis has often been criticized in the literature.

[6]Howson says that probabilities are relative to background information and that, in judging whether or not E confirms H, we should use as background information the current background information minus E. But to understand this we need to understand what distinguishes "background" from other information, what it means to subtract a proposition from a given body of background information (Chihara 1987), and what it means for p to be a person's probability function relative to some background information other than the person's actual background information. Howson's thought with respect to the last point is apparently that p represents my probability function with respect to background information K iff p is the probability function I would have if K were my background information (1991, 548). This last point is what makes Howson's analysis a counterfactual one.

One criticism is that it is not always clear precisely what probability function one would have if one did not know something one does know (Glymour 1980, 87ff.; Garber 1983, 103; Eells 1985, 286; etc.) However, for the purposes of an analysis of qualitative confirmation we do not need to identify a precise probability function; even if it is not clear what precise probability function one would have if one did not know E, it might still be clear whether any such probability function would satisfy the condition $p(H|E) > p(H)$. Furthermore, it is not necessary to insist that even this always be clear, since one might say (with Howson 1991, 449f.) that confirmation is sometimes indeterminate.

A slightly more formidable objection is that there are cases in which we judge that E does confirm H but in which it is arguably the case that, were we not to know E, we would not have considered H, so that we would not have a probability for H, and so would not have a probability function for which $p(H|E) > p(H)$ (Eells 1985, 286; Earman 1992, 123). This sort of objection can be met by a small modification to the counterfactual analysis. A counterfactual analyst can say that E confirms H for X at t iff, were X not to know E *but to have probabilities for H and E* then, if p is the probability function X would have, $p(H|E) > p(H)$. (The emphasized clause is new.)

However, this revised counterfactual analysis is still open to criticism. There are contexts in which, were we not to know E, then even if we had probabilities for H and E still we might not properly appreciate the relationship between E and H. Thus it can be the case that we judge E to confirm H but, were we not to know E, we would have a probability function p such that $p(H|E) \leq p(H)$. The following is an example of this kind.

Suppose Mr. Schreiber is the author of novels that are popular (P) though it is important to him that he is making important contributions to literature (I). Schreiber basks in his success, taking his popularity to be evidence of the importance of his work; that is, he takes P to confirm I. However, he is well aware that the many aspiring serious novelists whose work is unpopular tend to rationalize their failure by supposing that the public taste is so depraved that nothing of true value can be popular. Schreiber thinks this reaction to unpopularity is unjustified and due merely to an inability to admit that one's work lacks merit. However, we can suppose that if Schreiber did not know of his own work's popularity, he too would share this opinion. We can even suppose that Schreiber, aware of his own foibles, is aware of this fact. Then it is true, and Schreiber knows it is true, that were he not to know P, he would have a probability function p such that $p(I|P) \leq p(I)$. However, Schreiber now thinks that his judgment in this counterfactual situation would be irrational. Thus he now judges that P confirms I though the counterfactual analysis implies the contrary.

Van Fraassen (1988, 155) argues that counterfactual analysis may go similarly astray in the Mercury perihelion example.

How could the counterfactual analysis be modified to avoid such counterexamples? We would have to say that the probability function to use for the purpose of evaluating whether E confirms H for X at t is one X would have if X did not know E, X gave probabilities to H and E, *and X had probabilities that X judges to be rational in that situation.* This emendation, however, will achieve the desired result only if, in cases in which X takes E to confirm H, X judges that rationality requires him in the relevant counterfactual situation to have a probability function such that $p(H|E) > p(H)$.

What norms of rationality must X endorse to make such a judgment? For any propositions H and E there are coherent probability functions p such that $p(E) < 1$ and $p(H|E) \leq p(H)$.[7] Thus if one did not know E one could still always have $p(H|E) \leq p(H)$ without violating coherence. Hence to judge that rationality requires $p(H|E) > p(H)$ when E is not known, X must endorse norms of rationality stronger than coherence.

Adding that other norm endorsed by subjectivists, namely conditionalization, does not change the situation. Schreiber, for example, need not have violated conditionalization in the counterfactual situation in which he does not know P. More fundamentally, even if a probability distribution has been reached via a violation of conditionalization, the conditionalization principle does not imply that it is now irrational to have that probability distribution; on the contrary, any correction of past violations of conditionalization will themselves constitute further violations of conditionalization.

Let me now review the argument to this point. We saw that subjectivist Bayesian analyses of confirmation are motivated by the desire to avoid commitment to any norms for probability judgment beyond coherence and conditionalization. I suggested that these analyses are most charitably understood as trying to analyze what it is for a person to take E to be such that, if it were evidence, it would be evidence for H. We found that the non-counterfactual subjective Bayesian analyses of this concept are irremediably flawed. We also found that the counterfactual analysis can be turned into an adequate analysis of subjective confirmation only if people endorse norms of rationality that go far beyond the norms of coherence and conditionalization endorsed by subjective Bayesianism. I conclude that subjective Bayesianism can provide a tenable analysis of the psychological concept it is trying to analyze only on the assumption that people endorse norms of rationality much stronger than those endorsed by subjective Bayesianism. In particular, since subjective Bayesians them-

[7]A trivial proof is to take a probability function in which $p(H) = 1$. The inequality can also be satisfied if $p(H) < 1$ provided that H neither entails nor is entailed by E.

selves make substantive first-person judgments of confirmation, they must allow that they themselves accept such norms.

Now the reason why subjective Bayesians made confirmation relative to persons and times was that the norms of rationality they endorsed were too weak to provide any analysis of objective confirmation such as Carnap offered. However, now that we see that stronger norms must be conceded anyway, this motivation for the subjective approach has been undermined. If we are going to permit ourselves to assume stronger norms of rationality, we may be able to offer a satisfactory analysis of confirmation in the objective sense, i.e., an analysis of what it is for E to *be* evidence for H, rather than of what it is for someone to *take* it to be so. We could then say that X judges E to be evidence for H iff X judges that this objective relation holds. This would have the advantage of freeing the analysis of confirmation from reliance on counterfactual conditionals.

3. Proposal. In this section I will offer an analysis of confirmation in the objective sense. I will argue that this analysis avoids the weaknesses of the preceding analyses—not only the weaknesses that I have already identified in those analyses but also others that I will point out below.

3.1. Evidence. Unlike Carnap and the subjectivist Bayesians, my analysis will require E to be a piece of evidence in order for E to confirm anything. One reason for this is that I want 'E confirms H' to correspond to our ordinary notion of E being evidence for H and this requires E to be evidence. I have another reason too, which I will explain in Section 3.4; I will also note there that Carnap's motivation for not requiring E to be evidence is undermined by my analysis.

I propose that E is evidence iff E is known directly by experience. Before I elucidate this proposal, let me note some immediate applications of it.

A proposition that is not known to be true is not evidence and hence is not evidence for anything. So, for example, the proposition that the Moon is made of cheese is not evidence that Io is made of cheese, because we do not know that the moon is made of cheese (in fact, we know this is false).

Even if a proposition is known to be true, if this knowledge is not directly based on experience then E is not evidence and hence is not evidence for anything. For example, let E denote that a substance taken from a certain jar dissolved when placed in water, and suppose we know E directly from experience. Let H be the proposition that the substance that was in the jar is soluble and suppose we know H by inferring it from E. Then we know both H and E from experience, but E is evidence for H and not vice versa, since only E is known *directly* by experience.[8]

[8]This example is mishandled by Achinstein's (1983) definition of confirmation. For Achinstein says that H is (potential and veridical) evidence that E if (i) E and H are true, (ii) H

On my account, the notion of evidence depends on that of knowledge and it may be asked what knowledge is. I intend this concept to be understood in its usual sense. Exactly what that sense is is a question on which epistemologists have spilled much ink. The question usually discussed in this literature is what it means for a person X to know a proposition P. It is relatively uncontroversial that X knows P only if P is true and X believes P. It is also relatively uncontroversial that these conditions are not sufficient; the main point of controversy is the exact nature of the condition or conditions that must be added in order to have a sufficient condition for knowledge. Some say that X must have a justification of a suitable kind, though what a justification is and what would be a suitable kind of justification are also controversial. Others say that X's belief in P must have been caused in a suitable way, without necessarily supposing that such causation constitutes a justification. Fortunately I do not need to settle these disputed points here; nothing that I have to say about evidence depends on what exactly needs to be added to truth and belief to have a correct analysis of knowledge.

A question that is more important for my purposes is what it means to say that *X believes P*. Traditionally it was supposed that X can know P only if X is certain of P. Epistemologists today, having recognized that few of the things we say we know are things of which we can be absolutely certain, have generally jettisoned this requirement. However, if X is not certain of P then X would be rational to bet that P is false if the odds against P are high enough, and it is quite paradoxical to say that betting on a proposition that one knows is false can be rational.

I suggest that we reconcile these conflicting considerations by taking certainty to be a relative notion. What is really required for knowledge is not absolute certainty (a probability of 1) but rather a probability that is high enough that it can be treated as certainty for all purposes that are under consideration. To give certainty in this sense a name I will call it *practical certainty.*

It might seem that, in order to know whether P can be treated as certainly true for all purposes under consideration, we must determine the probability of P and thereby determine the rationality of acting as if P is true for all purposes under consideration. However, if such a calculation is required then P is not practically certain in the sense I intend. A practically certain proposition is one in which it is clear, without doing any

---

does not entail E, (iii) there is an explanatory connection between the truth of H and E, and (iv) this explanatory connection is probable given that H and E are true. All these clauses can be supposed to be satisfied in my example. In particular, (ii) is satisfied because H does not entail that the substance will be placed in water, which is part of what E asserts. Thus Achinstein deems H to be (potential and veridical) evidence that E, though in fact it is neither because it is not evidence.

calculations, that no such calculation is necessary. Indeed, it will normally be irrational to perform any such calculations for a practically certain proposition (cf. Maher 1993, Sec. 1.2). Practical certainties are thus treated just like absolute certainties for the purposes at hand.

Example: If our purpose is to evaluate the probability of the General Theory of Relativity (G) in the light of the available evidence, the calculation would be considerably simplified by supposing that we are certain of the advance of the perihelion of Mercury (A), along with other similar evidence. Furthermore, I suppose that someone familiar with the field does not need to do a calculation to see that treating A as certain will not introduce any significant error. If this supposition is correct, then in such a context A counts as practically certain. On the other hand, if for some reason we are considering the question of whether it would be rational to accept a million-to-one bet against A, then we should not proceed by assuming that A is true but should instead evaluate the probability of A; it is then no longer true that we are practically certain of A and, for that reason, we would not say that we know A (not even if our evaluation shows that we should not bet against A).

There are situations in which, although we have learned something from experience, there is no proposition of which we are practically certain that expresses what we have learned. In requiring E to be practically certain in order to count as evidence I am not denying the existence of such situations. Such situations are ones in which the evidence is not propositional (Jeffrey 1965, Ch. 11). What I am saying is that *if* proposition E is truly evidence then E must be practically certain.

Not all propositions that are known to be true are evidence. On the account I am proposing, evidence is known *directly* by experience. But what does it mean for a proposition to be known directly by experience, rather than indirectly? The traditional conception was that a proposition is known directly by experience if knowledge of it is not based on inference from any other proposition. However, we say that A is evidence and our knowledge of it plainly depends on inference from other propositions. Originally A was inferred from the results of various particular observations by astronomers; now, for most of us, it is inferred from assertions of authorities.

The considerations just cited might suggest that evidence can be known on the basis of an inference from other propositions. However, the particular observations from which A was inferred are surely also evidence, and it seems at best peculiar to say, in the same breath, that both the particular observations and the phenomenal law inferred from them are evidence.

These superficially conflicting tendencies all find a natural place if we understand claims about directness as relative to a set S of propositions that are taken as relevant. A proposition E will then be said to be known directly by experience iff it is known by experience and that knowledge is not based on an inference from any proposition in S. Usually, in discussions of G, the particular observations from which A was inferred are not being considered, hence A counts as known directly by experience and so can count as evidence. Conversely, if the particular observations were being considered, we would call them the evidence and regard A as an inferred proposition, not as evidence.

Suppose, though, that S includes both A and G. Our confidence in A depends largely on observation, but it should also be increased by the fact that A follows from G and there is other evidence for G (such as the bending of light in a gravitational field). Does this mean that knowledge of A is based on an inference from a proposition in S, in which case A does not count as known directly by experience? No, it does not mean that. While our other reasons for believing G make A somewhat more probable than it would otherwise be, I take it that A would still be practically certain in the absence of any reason to believe G. History bears out this supposition, since A was in fact practically certain before the formulation of G, at which time it was an anomaly for the then-accepted theory of gravitation.

A final point: If evidence is knowledge obtained directly by experience, and if knowledge is a relation between a person and a proposition (as the analyses of epistemologists usually assume), then what counts as evidence will be different for different people. While we sometimes talk this way, there is also a notion of evidence in which it is regarded as a property of a community rather than of individual persons. For example, we can say that A was part of the evidence possessed by the physics community in the late 19th century. This does not mean that every physicist in that community knew A, much less that they all knew A directly by experience. It does not even mean that *most* physicists did so. What it means is that someone had learned A directly by experience and this knowledge had been made available to the community along with an appropriate justification. This communal notion of evidence is the usual one in science; when scientists claim that "the evidence" supports hypothesis H they are making a claim about the evidence available to the community at that time, not merely about the evidence that they know.

Although E may be evidence for a community without being evidence for every individual in the community, it is normally the case that if E is evidence for the community then it is evidence for every individual in the community who knows E. For suppose Y has learned of E from the report of some other individual X, who in turn justifies E from experience; then although Y has inferred E from X's report, Y will still count as knowing E directly by experience provided propositions about other people's reports are not under consideration, as is commonly the case; thus E will count as evidence for Y as well as for X.

To sum up: A necessary condition for E to confirm anything is that E be evidence. The notion of evidence is relative to an individual or a community; I have not specified this explicitly but have left it to be determined by the context. For an individual X, E is evidence iff X knows E directly by experience. The knowledge in question requires practical certainty of E, which is relative to the ends in view. Such knowledge counts as direct from experience if it is not inferred from any of the propositions under consideration. For a community C, E is evidence iff E has been evidence for some individual who made this information available to C along with an appropriate justification.

3.2. Probability. I now turn to the question of what it is for some evidence E to confirm, or be evidence for, H. The basic idea I want to articulate and defend is that evidence E is evidence for H iff E makes it rational to be more confident that H is true. This was Carnap's basic idea, although my way of articulating it will differ from his.

As Carnap recognized, whether or not E makes it rational to be more confident of H depends on what other evidence is assumed to be available. For example, we saw earlier that (assuming certain other facts about the solar system) A makes it rational to be more confident of G if S is assumed to be a piece of available evidence, while A makes it rational to be *less* confident of G if $\bar{S}$ is assumed to be part of the available evidence. I will call this other evidence *background evidence.*

I will view the background evidence as a set of pieces of evidence. Elements of the background evidence will sometimes be representable by a proposition that satisfies the conditions for evidence given in the preceding section. However, as noted there, what we learn from experience is not necessarily representable by a proposition we can formulate, so some items of background evidence may not be able to be expressed as a proposition. Nevertheless, we can represent non-propositional evidence by a variable that is treated for most purposes the same as a proposition (Pearl 1990) and that is the procedure I will follow here. The basic idea can then be stated more precisely as follows: E is evidence for H given background evidence $\mathcal{B}$ iff it is rational to be more confident of H if one's total evidence is $\mathcal{B} \cup \{E\}$ than if it is $\mathcal{B}$ alone.

In common with both Carnap and subjective Bayesians, I subscribe to the view that rational degrees of confidence are representable by a numerical measure that satisfies the axioms of probability. (For a defense of this, see Maher 1993.) Therefore, instead of talking about degrees of confidence that it is rational to have, I will talk about rationally permissible probability functions. Now, what probability functions it is rational to have in a given situation clearly depends on what evidence is available in that situation. I will therefore use the notation $\mathcal{R}(\mathcal{B})$ to denote the set of probability functions that are rationally permissible when the total evidence is $\mathcal{B}$. By this I mean that $p \in \mathcal{R}(\mathcal{B})$ iff there is a possible situation in which some person's total evidence is $\mathcal{B}$, that person has probability function p, and that person does not violate any norm of rationality.

I do not assume that $\mathcal{R}(\mathcal{B})$ contains only one probability function; perhaps different people can have different probabilities given the same evidence without either being irrational. Nor do I assume that what probability functions a person can rationally have is determined solely by the person's evidence; it might, for example, depend also on what prior probabilities the person had. Thus I am not assuming that if $p \in \mathcal{R}(\mathcal{B})$ then any person with total evidence $\mathcal{B}$ could rationally have probability function p. I am merely saying that for every $p \in \mathcal{R}(\mathcal{B})$ there is a possible situation in which a person with total evidence $\mathcal{B}$ rationally has probability function p.

Also in common with both Carnap and subjective Bayesians, I endorse the principle of conditionalization under conditions that we can here assume to be satisfied (Maher 1993, Ch. 5). Thus for a person with probability function p, learning E makes it rational to be more confident of H iff $p(H|E) > p(H)$.

Putting these points together leads to the following analysis of confirmation:

> E confirms H relative to $\mathcal{B}$ iff (i) E is evidence and (ii) for all $p \in \mathcal{R}(\mathcal{B})$, $p(H|E) > p(H)$.

I think this analysis is correct as far as it goes, but it does have a serious limitation. We often want to say that several pieces of evidence jointly confirm H, without implying that any of the individual pieces of evidence separately does so. For example, we might want to say that A and S confirm G jointly but not separately. The analysis just given, however, only covers the case where a single piece of evidence is said to confirm something. We might try saying that $E_1, \ldots, E_n$ jointly confirm H iff their conjunction does so. However, if $E_1, \ldots, E_n$ are evidence their conjunction would normally be inferred from them, in which case the conjunction has not been made certain *directly* by experience and so does not count as evidence. (Nor can we avoid this by saying that $E_1, \ldots, E_n$ are not being considered here, since the statement we are trying to analyze mentions them.) In any case, the interpretation of joint confirmation as confirmation by a conjunction strikes me as unnatural. What I propose instead is that we think of confirmation in the general case as a relation between a *set* of pieces of evidence and a hypothesis. This suggests the following analysis:

> $\{E_1, \ldots, E_n\}$ confirms H relative to $\mathcal{B}$ iff (i) $E_1, \ldots, E_n$ are all evidence and (ii) for all $p \in \mathcal{R}(\mathcal{B})$, $p(H|E_1 \ldots E_n) > p(H)$.

This last analysis is the one I want to defend in the general case. However, I will follow ordinary usage in not distinguishing between E and $\{E\}$, so that when a single piece of evidence is under consideration the earlier analysis can be used.

3.3. Constraints. According to my analysis, judgements about what confirms what, given $\mathcal{B}$, depend in significant part on what probability functions are in $\mathcal{R}(\mathcal{B})$. So I will say what I can about the contents of $\mathcal{R}(\mathcal{B})$.

First, the axioms of probability limit what counts as a probability function, and thus all elements of $\mathcal{R}(\mathcal{B})$ must satisfy those axioms. As noted earlier, this constraint is often called *coherence*.

Second, every $p \in \mathcal{R}(\mathcal{B})$ needs to be consistent with $\mathcal{B}$. Since evidence must be practically certain, what this condition amounts to is that $p(E)$ is practically 1 for all $E \in \mathcal{B}$ and $p \in \mathcal{R}(\mathcal{B})$.

What about norms governing how probabilities should be revised as new evidence is acquired? Since evidence must be practically certain, the only such norm that need concern us is conditionalization. Now any coherent probability function that gives every element of $\mathcal{B}$ probability 1 is the result of conditioning some probability function on every element of $\mathcal{B}$; consequently, the principle of conditionalization does not exclude from $\mathcal{R}(\mathcal{B})$ any probability function that is not excluded by the preceding two conditions.

The constraints I have mentioned so far are relatively uncontroversial, being endorsed by subjective Bayesians and logical probability theorists alike. However, subjective Bayesians do not endorse any stronger constraints than those mentioned so far, and here my analysis requires me to part company with them. For if $\mathcal{R}(\mathcal{B})$ contained all the probability functions that satisfy the constraints mentioned so far, it would be so large and diverse that few judgments of confirmation would be deemed true by my analysis. But we have seen that subjective Bayesians do not succeed in analyzing even their subjective notion of confirmation without endorsing stronger norms, so my need to appeal to such stronger norms is not a disadvantage of my analysis as compared with a subjective Bayesian one.

However, I don't wish to rest my case merely on a *tu quoque* argument, so I will say what I can in support of the reality of these additional norms of rationality.

There are probability functions that satisfy the rationality conditions mentioned so far but which most of us are inclined to call absurd, unsupported, paranoid, or the like. These epithets are not mere descriptions but embody a normative evaluation of the negative kind. Following Gibbard (1990), I take it that to call something irrational is, roughly, to express one's acceptance of norms that do not permit it. I conclude that there are opinions, which violate none of the norms mentioned earlier, but which are irrational.

To illustrate: It is compatible with both coherence and conditionalization to regard all possible worlds consistent with our evidence as having the same probability. However, as Carnap (1950, 565) pointed out, this way of assigning probabilities implies that no evidence is relevant to what we have not yet experienced. For example, someone who assigned probabilities in this way and who knew that millions of ravens had been observed to date, all of them black, would still think that the next raven was no more likely to be black than any other color. Such a person would be irrational, despite satisfying both coherence and conditionalization.

Carnap (1950, 1952, 1971b), Rosenkrantz (1977), Salmon (1988), and others, have tried to articulate general norms of rationality governing probability functions that go beyond coherence and conditionalization. This is not the place to review the adequacy of such proposals, but I will say that I am inclined to concede to the subjectivist that no such proposal provides a general and substantive norm of rationality that is as precise and compelling as coherence and conditionalization. However, this concession is not a good reason to deny that we endorse any norms beyond coherence and conditionalization. Even in the absence of a general, precise, and compelling codification of such norms, we clearly do have a patchwork of piecemeal normative judgments about particular situations, judgments which do not follow from the norms endorsed by subjective Bayesianism.

Since we lack a complete theory of rationally permissible probability functions, it will in some cases be controversial whether or not all $p \in \mathcal{R}(\mathcal{B})$ are such that $p(H|E) > p(H)$. I think that in such cases it will also be controversial whether or not E confirms H, and hence I think this consequence of my analysis is just what it should be. There are certainly many cases in the history of science in which it was disputed whether or not a particular piece of evidence supported a particular theory. My analysis does not aim to resolve all such disputes but rather to identify what claim is being disputed.

3.4. Context. The analysis of confirmation that I have given makes it relative to background evidence $\mathcal{B}$. In this respect, my analysis as I have presented it so far is similar to Carnap's. However, when we make judgments of confirmation we do not usually specify any background evidence. For example, when the advance of the perihelion of Mercury (A) is said to confirm the General Theory of Relativity (G), background evidence is normally not mentioned. Such assertions of confirmation have the surface grammar of what Carnap called "absolute confirmation," but so interpreted they would come out false; A does not confirm G relative to no background evidence. It is a defect of Carnap's analysis that it does not provide any more satisfactory way to interpret such claims.

I propose that the background evidence is usually fixed sufficiently well by the context and that is why it is normally not explicitly stated. This is not an unusual situation; there are many sentences whose meaning is fixed in part by the context in which they are uttered. For example, if I walk into a meeting of my department and say "Steve is coming," my hearers will understand that I am saying my colleague Steve is on his way to the meeting. On the other hand, if we have been discussing the dinner organized for a visiting speaker that evening and I say "Steve is coming," my hearers will understand that I am saying that Steve will go to the dinner. So we see that what the sentence "Steve is coming" expresses is that Steve is going towards some reference point p, though p is not explicitly specified but rather is taken to be fixed by the context. Of course, we can explicitly indicate a reference point if we wish, and will do so if the reference point would otherwise be unclear or inappropriate; thus if tonight's dinner has not been the subject of the immediately preceding conversation, I will say "Steve is coming to dinner tonight," thereby specifying the reference point.

I am proposing that 'E confirms H' is similar to this; what it states is that E confirms H relative to background evidence $\mathcal{B}$, but $\mathcal{B}$ is usually taken to be fixed to sufficient accuracy by the context and so is normally not explicitly indicated. We can if we want explicitly indicate what the background evidence is, and in some contexts we will need to specify it in order to say what we want to say; but usually no such specification is needed or provided.

How is the background evidence determined by context? A natural suggestion, and one that parallels Howson's formulation of the counterfactual analysis, is to say that, if we are considering whether E confirms H, the background evidence is all the evidence available in the context other than E. However, this suggestion leads to erroneous judgments of confirmation, as the following examples show.[9]

Suppose detectives at the scene of a crime have found two fingerprints that correspond to those of Jones. Let F_1 denote the finding of the first fingerprint and F_2 the finding of the second set, and let $\mathcal{B}$ denote the total available evidence other than F_1 and F_2. We can suppose that, given $\mathcal{B}$, one fingerprint is just as good evidence as two that Jones was at the scene of the crime. That is, we can suppose that, for all $p \in \mathcal{R}(\mathcal{B})$, $p(J|F_1) = p(J|F_2) = p(J|F_1F_2) > p(J)$. Then the total evidence other than F_1 is $\mathcal{B} \cup$

[9]Hence the following examples are also counterexamples to the version of the counterfactual analysis defended by Howson. I did not press this point earlier because the counterfactual analysis could be modified to use the alternative rule for determining background evidence that I will propose.

$\{F_2\}$ and, for all $p \in \mathcal{R}(\mathcal{B} \cup \{F_2\})$, $p(J|F_1) = p(J)$, so that F_1 does not confirm J according to the present proposal. An exactly parallel analysis shows that F_2 also does not confirm J. Thus neither fingerprint counts as confirming J because it adds nothing to the case based on the other one. This is wrong.

In the fingerprint example the two pieces of evidence support the hypothesis independently. Another kind of counterexample to the proposed way of fixing background evidence arises when we have pieces of evidence that only support the hypothesis together, not separately. The evidence A and S for G (the General Theory of Relativity), discussed earlier, is a case of this kind. If both A and S are part of our evidence, then the proposed way of fixing background evidence has the result that A and S both count as evidence for G. However, to say that each of these propositions is evidence for G is to imply that we have here two separate supports for G when in fact there is only one support; that support is more fully described as the set $\{A, S\}$.

So when we are considering whether E confirms H, the background evidence should not in general be identified with the total available evidence other than E. What I propose instead is this: If propositions $E_1, \ldots, E_n$ are all the propositions whose relevance to H is under discussion or otherwise salient in the context then, *celeris paribus,* $\mathcal{B}$ is the total evidence available in that context other than $E_1, \ldots, E_n$. Putting aside the *ceteris paribus* clause, this rule agrees with the one I have just rejected iff there is only one proposition whose relevance to H is under discussion or otherwise salient in the context.

The reason for the *ceteris paribus* clause is that other considerations sometimes influence the content of the background evidence. For example, within certain limits the background evidence may be fixed by the principle that it is whatever it needs to be to make the assertions in that context true. We see analogs of this in other context-dependent statements. For example, if I say "Steve is coming because he wants to talk with the speaker" the reference point is fixed as the dinner, even though I have not stated so explicitly and even if the dinner had not been the topic of discussion. Similarly, if someone says "E is (not) evidence for H" then, *ceteris paribus,* that limits the background evidence to something relative to which E is (not) evidence for H. Lewis (1979) has called principles of this kind *rules of accommodation.*

How do we identify "the total evidence available other than $E_1, \ldots, E_n$"? The total evidence available can be regarded as a set of pieces of evidence, each of which is known directly by experience (in the sense explained in Section 3.1). As mentioned earlier, I am not assuming that these items are all propositions, but I am assuming that they can be treated as if they were. So if $\mathcal{T}$ represents the total evidence available then the total evidence available other than $E_1, \ldots, E_n$ is simply the set consisting of all elements of $\mathcal{T}$ other than $E_1, \ldots, E_n$, i.e., the set $\mathcal{T} - \{E_1, \ldots, E_n\}$.

Some readers will be familiar with the similar-sounding problem of identifying the corpus of beliefs that results from removing a belief from the corpus. The latter problem is made difficult by the fact that belief $\mathcal{B}$ in corpus $\mathcal{K}$ may be entailed by others beliefs in $\mathcal{K}$, so removing $\mathcal{B}$ leaves $\mathcal{K}$ not deductively closed, whence the question arises as to what else should be removed—a question to which there seems to be no natural answer (Gärdenfors 1988, Section 3.4). This difficulty does not arise for me because I am discussing sets of pieces of evidence, not corpora of beliefs.[10]

Earlier I gave one reason for requiring E to be evidence in order for E to confirm anything: This is necessary for the notion of E confirming H to correspond to the ordinary concept of E being evidence for H. I can now state another reason: If I did not require E to be evidence, my analysis of confirmation would need to be complicated in a way that would serve no good purpose. This is because, so long as E was assumed to be evidence, I only needed to say how $\mathcal{B}$ is fixed by context for contexts in which E is part of the available evidence. If confirmation did not require E to be evidence, my account of confirmation would require me to extend my account of how $\mathcal{B}$ is fixed by context to cover contexts in which E is not part of the available evidence, including cases where the available evidence is inconsistent with E. I don't see any obviously right way of doing this. Indeed, I don't even see any natural criterion for deciding whether a particular way of fixing $\mathcal{B}$ in such cases is right or not.

In discussing Carnap I explained his motivation for excluding from the analysis of confirmation the requirement that E be evidence; it was that whether E confirms H should not depend on such factual questions as whether E is true or known to be true. However, if what I have said is at all on the right track, determination of the truth value of 'E confirms H' normally requires knowledge of the context in which this sentence was uttered and the evidence available in that context. Thus, even if I were to drop from my analysis of confirmation the requirement that E be evidence, one would still need factual knowledge, and in particular knowledge of

[10]The problem of subtracting a belief from a corpus was argued by Chihara (1987) to be a problem for Howson's version of the counterfactual analysis. This is because Howson said that, when judging whether or not E confirms H, we should use as background information the current background information minus E; Chihara's question was how E is to be subtracted from the background information. Howson's response (1991) involved the claim that people have axiomatized their background information and, in cases where E is one of the axioms, we simply remove it. It isn't clear to me what Howson's "background information" is, but whatever it is, it seems implausible that people have any unique axiomatization of it. Howson's proposal also does not tell us what to do when E is not one of the axioms, as he concedes. Howson could avoid these unsatisfactory consequences if he replaced 'background information' by 'evidence' and limited claims of confirmation to propositions that are evidence.

the available evidence, in order to determine what confirms what. This negates Carnap's motivation for not requiring E to be evidence.

3.5 Examples. I will now illustrate the application of my analysis by applying it to two examples that I discussed earlier.

Suppose we are considering whether or not A (the observed advance of the perihelion of Mercury) confirms G (the General Theory of Relativity). I am supposing that the status of other pieces of evidence is not also under discussion or otherwise salient. Then if $\mathcal{T}$ is the total available evidence, the background evidence will, *ceteris paribus,* be $\mathcal{T} - \{A\}$. If $\mathcal{T}$ is what we now take it to be (including S) then it is relatively uncontroversial that, for all $p \in \mathcal{R}(\mathcal{T} - \{A\})$, $p(G|A) > p(G)$. Hence A does confirm G.

Similarly, if we are considering whether S (the sphericity of the sun) confirms G, and all other evidence (including A) is being taken for granted, then S will count as confirming G.

Now suppose that both A and S are being considered with regard to their evidential relevance to G, or are otherwise salient in the context, and these are the only pieces of evidence whose relevance to G is under discussion or otherwise salient. Then, *ceteris paribus,* the background evidence will be $\mathcal{T} - \{A, S\}$. Assuming as before that $\mathcal{T}$ is what we now take it to be, it is relatively uncontroversial that, for all $p \in \mathcal{R}(\mathcal{T} - \{A, S\})$, $p(G|AS) > p(G)$. Hence $\{A, S\}$ confirms G. However, it is arguably the case that, for at least some $p \in \mathcal{R}(\mathcal{T} - \{A, S\})$, $p(G|A) \leq p(G)$ and $p(G|S) \leq p(G)$. Hence neither A nor S by itself confirms G.

The results of the last three paragraphs will seem contradictory if one forgets that an assertion of confirmation makes a claim that is relative to context. Different contexts fix the background evidence in different ways and so what might seem to be the same assertion (e.g., 'A confirms G') can have different truth values in different contexts. This is perfectly normal for statements whose meaning is determined in part by context. The statement "Steve is coming" is likewise true in some contexts and false in others.

Note that, on my analysis, while it can be correct to say A is evidence for G, and it can be correct to say that S is evidence for G, it cannot be correct to say both these things. The reason is that the background evidence needs to be different for each statement to be true and, in one statement, the background evidence cannot be two different things. This also corresponds to what happens with other indexical statements. It can be true to say that Steve is coming and it can be true to deny it, but it cannot be true that Steve is both coming and not coming, for there is no reference point that can make both statements true.

I turn now to the fingerprint example of Section 3.4. In this context the evidential status of both F_1 and F_2 is likely to be under discussion, or at least both propositions are apt to be salient since they are so similar. Assuming that no other evidence is also under discussion or otherwise salient, the background evidence is, *ceteris paribus,* $\mathcal{T} - \{F_1, F_2\}$. For the sake of the example I suppose that, for all $p \in \mathcal{R}(\mathcal{T} - \{F_1, F_2\})$, $p(J|F_1) = p(J|F_2) = p(J|F_1F_2) > p(J)$. Thus F_1 confirms J and so does F_2, as does the combination $\{F_1, F_2\}$. What is true is merely that $\{F_1, F_2\}$ is not stronger evidence for J than F_1 and F_2 are separately.

4. Achinstein. All the analyses of confirmation considered in this essay are attempts to explicate the *relevance concept* of confirmation, according to which E is evidence for H iff E makes H more probable. Achinstein (1983) maintains that this relevance concept of confirmation is not the one that is ordinarily used in science and everyday life. He has defended this position with examples in which, he claims, our ordinary judgments about what is evidence for what diverge from the relevance concept. He has also offered his own analysis of confirmation according to which E can be evidence for H without E making H more probable (cf. note 8). Bar-Hillel and Margalit (1979) and Kronz (1992) have disputed Achinstein's examples and his conclusion but Achinstein (1981, 1992) has rejected these criticisms. To deal fully with the issues raised by Achinstein is beyond the scope of this paper, but I will comment on some of Achinstein's putative counterexamples to the relevance concept.

Achinstein says: "When Mark Spitz [an Olympic swimmer] goes swimming he increases the probability that he will drown; but the fact that he is swimming is not evidence that he will drown" (1983, 152).[11] Now if Spitz might be going swimming in unsafe circumstances, such as a flooded river, even Achinstein would agree that this is evidence that he will drown. So let us assume that the evidential proposition is that Spitz is going for a training swim in the usual way in a normal pool at some particular time; let this proposition be M. Let D denote that Spitz will drown on this occasion. In assessing whether or not M is evidence for D we should use background evidence $\mathcal{B}$ that does not include M but does include our other evidence, such as that Spitz is an outstandingly competent swimmer, that drowning is in general more likely when a person is in water than otherwise, but also that there are no known cases of a swimmer of Spitz's competence drowning on a training swim. I allow that, for all $p \in \mathcal{R}(\mathcal{B})$, $p(D|M)$ is greater than $p(D)$, even though it is only infinitesimally so. Since it is greater, on my account M does count as evidence for D. Achinstein thinks this is obviously wrong. I agree that it would normally be inappro-

[11]As we see from this example, Achinstein talks of evidence *that* a hypothesis is true, whereas I talk of evidence *for* a hypothesis. For Achinstein's objections to be relevant to the account of confirmation defended here, this difference needs to be regarded as merely stylistic and I will so regard it.

priate to say that M is evidence for D, but it does not follow that this statement is false.

It is essential to Achinstein's example that $p(D|M)$ and $p(D)$ are known to be very small. That being the case, the question of whether or not D is true, or even what its probability is, will not normally arise. But the statement that M is evidence for D would normally only be a relevant contribution to a conversation if there were some question about whether or not D was true. Since a statement needs to be relevant to be appropriate (Grice 1989), it follows that it will normally be inappropriate to say that M is evidence for D.

If this account is right then, if we can find a context in which it is relevant to say that M is evidence for D, it ought to be appropriate to make this statement in that context. Here is such a context: A group of philosophers marooned on a desert island are passing the time with idle intellectual games. One day one of them proposes the game in which the players have to give examples of confirming evidence, the winner being the one whose example involves the weakest confirmation. If one of the players defines M and D as we have done here and says that M is evidence for D, then I think he has made an appropriate statement; what he has said is true and he should win the game unless another player can come up with an even weaker example of confirmation.[12]

Achinstein follows the Spitz example with this one: "When I walk across the street I increase the probability that I will be hit by a 1970 Cadillac; but the fact that I am walking across the street is not evidence that I will be hit by a 1970 Cadillac" (1983, 152). The chance of being hit by a 1970 Cadillac while crossing the street is normally about as small as that of Mark Spitz drowning while swimming, so the statement "My walking across the street is evidence that I will be hit by a 1970 Cadillac" will normally be inappropriate for the same reason as in the Spitz example. But there is an additional source of inappropriateness in the present example: For the statement to be fully relevant, there normally must be something salient about 1970 Cadillacs, such as that I am more likely to be hit by one of them than by other cars. This salience of 1970 Cadillacs is what Grice (1989) calls a conversational implicature; its falsity is an additional reason why it would normally be inappropriate to say that my walking across the street is evidence that I will be hit by a 1970 Cadillac. However, this statement would be a perfectly appropriate one in the game of our desert island philosophers.

[12] I expect that Achinstein would allow that M is slight evidence that D *may* be true but deny that it is even slight evidence that D. I don't think this distinction has much intuitive support, so I will not here attempt to show that it is mistaken. My present purpose is merely to show that intuitive judgments that might seem to conflict with my analysis can actually be done justice in a way that is consistent with that analysis.

The third and last example I will discuss raises different issues. Achinstein writes:

> Let $\mathcal{B}$ be the background information that there is a lottery consisting of 1001 tickets, one of which will be drawn at random, and by Tuesday 1000 tickets have been sold, of which Alice owns 999. Let E be the information that by Wednesday 1001 lottery tickets have been sold, of which Alice owns 999, and no more tickets will be sold. Let H be the hypothesis that Alice will win. Now, I suggest, E is evidence that H in this case. The information that Alice owns 999 of the 1001 tickets sold and that no more tickets will be sold is evidence that Alice will win. (1983, 152, notation changed)

I agree that Achinstein's last sentence expresses a natural judgment, but that is because I think that in judging whether E confirms H in this case it is natural to exclude from the background evidence the proposition

> E_T: By Tuesday 1000 tickets had been sold, of which Alice owned 999.

The reason for excluding E_T is that it is highly salient, mainly because it is so similar to E. Furthermore, by the principle of accommodation (Section 3.4), the judgment that E is evidence for H fixes the background evidence as excluding E_T. Thus my own analysis can explain the plausibility of Achinstein's judgment that E confirms H.

This interpretation of the case does not respect Achinstein's statement that the background information includes E_T. However, when Achinstein states the judgment that he wants us to agree with, he does not include reference to this background information, and my point is that the plausibility of his judgment depends on that. We can see this by adding to Achinstein's judgment an explicit specification of the background evidence. The judgment becomes:

> Now, I suggest, E is evidence that H, given $\mathcal{B}$, in this case. The information that Alice owns 999 of the 1001 tickets sold and that no more tickets will be sold is evidence that Alice will win, given that on Tuesday Alice owned 999 tickets and only 1000 had then been sold.

This judgment does not seem to me to have any intuitive plausibility.[13] Thus I regard this lottery example as not a difficulty for my analysis of confirmation but rather a good example of how my analysis can consistently resolve what otherwise might seem conflicting intuitions about particular cases.

[13]Kronz (1992, 156) also argues that Achinstein's interpretation of this example is unnatural if we keep in mind what the background evidence is supposed to be. What I have added to his remarks is an explanation of why we tend to take the background information as excluding E_T.

I will end my consideration of Achinstein's examples at this point. There are other examples that I have not discussed, but I hope I have said enough to at least make plausible that Achinstein's examples do not show that my analysis of confirmation diverges from the concept of confirmation ordinarily used in science and everyday life.

REFERENCES

Achinstein, Peter (1978), "Concepts of Evidence", *Mind* 87: 22–45.

———. (1981), "On Evidence: A Reply to Bar-Hillel and Margalit", *Mind* 40: 108–112.

———. (1983), "Concepts of Evidence", in Peter Achinstein (ed.), *The Concept of Evidence.* Oxford: Oxford University Press, pp. 145–174. A modified version of Achinstein 1978.

———. (1992), "The Evidence Against Kronz", *Philosophical Studies* 67: 169–175.

Bar-Hillel, Maya and Margalit, Avishai (1979), "In Defense of the Classical Notion of Evidence", *Mind* 88: 576–583.

Carnap, Rudolf (1950), *Logical Foundations of Probability.* Chicago: University of Chicago Press. Second edition 1962.

———. (1952), *The Continuum of Inductive Methods.* Chicago: University of Chicago Press.

———. (1971a), "Inductive Logic and Rational Decisions", in Rudolf Carnap and Richard C. Jeffrey (eds.), *Studies in Inductive Logic and Probability,* Volume 1. Berkeley: University of California Press, pp. 5–31.

———. (1971b), "A Basic System of Inductive Logic, Part I", in Rudolf Carnap and Richard C. Jeffrey (eds.), *Studies in Inductive Logic and Probability,* Volume 1. Berkeley: University of California Press, pp. xx–xx.

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

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

Eells, Ellery (1985), "Problems of Old Evidence", *Pacific Philosophical Quarterly* 66: 283–302. Reprinted as "Bayesian Problems of Old Evidence" in Savage 1990.

Garber, Daniel (1983), "Old Evidence and Logical Omniscience in Bayesian Confirmation Theory", in John Earman (ed.), *Testing Scientific Theories.* Minneapolis: University of Minnesota Press.

Gärdenfors, Peter (1988), *Knowledge in Flux.* Cambridge, MA: MIT Press.

Gibbard, Allan (1990), *Wise Choices, Apt Feelings.* Cambridge, MA: Harvard University Press.

Glymour, Clark (1980), *Theory and Evidence.* Princeton: Princeton University Press.

Grice, Paul (1989), *Studies in the Way of Words.* Cambridge, MA: Harvard University Press.

Hempel, Carl G. (1945), "Studies in the Logic of Confirmation", *Mind* 54. Page references are to the reprint in Hempel (1965).

———. (1965), *Aspects of Scientific Explanation.* New York: The Free Press.

Horwich, Paul (1982), *Probability and Evidence.* Cambridge: Cambridge University Press.

Howson, Colin (1984), "Bayesianism and Support by Novel Facts", *British Journal for the Philosophy of Science* 35: 245–251.

———. (1985), "Some Recent Objections to the Bayesian Theory of Support", *British Journal for the Philosophy of Science,* xxxx.

———. (1991), "The 'Old Evidence' Problem", *British Journal for the Philosophy of Science* 42: 547–555.

Jeffrey, Richard C. (1965), *The Logic of Decision.* New York: McGraw-Hill. Second edition, University of Chicago Press 1983.

Kronz, Frederick M. (1992), "Carnap and Achinstein on Evidence", *Philosophical Studies* 67: 151–167.

Lewis, David (1979), "Scorekeeping in a Language Game", *Journal of Philosophical Logic* 8: 339–359.

Maher, Patrick (1993), *Betting on Theories.* Cambridge: Cambridge University Press.

Pearl, Judea (1990), "Jeffrey's Rule, Passage of Experience, and *Neo*-Bayesianism", in Henry

E. Jr. Kyburg, Ronald P. Loui and Greg N. Carlson (eds.), *Knowledge Representation and Defeasible Reasoning*. Dordrecht: Kluwer, pp. 245–265.

Rosenkrantz, Roger D. (1977), *Inference, Method and Decision*. Dordrecht: D. Reidel.

Salmon, Wesley C. (1988), "Dynamic Rationality: Propensity, Probability, and Credence", in James H. Fetzer (ed.), *Probability and Causality*. Dordrecht: Reidel, pp. 3–40.

Savage, C. Wade (ed.) (1990), *Scientific Theories*. Minneapolis: University of Minnesota Press.

Van Fraassen, Bas C. (1988), "The Problem of Old Evidence", in David F. Austin (ed.), *Philosophical Analysis*. Dordrecht: D. Reidel, pp. 153–165.

Will, C. M. (1986), *Was Einstein Right?* New York: Basic Books.