If we now let $j_n(\omega)$ stand for the number of heads in the first n trials of ω, the weak and strong form of the law of large numbers can be stated as follows:

WLLN The $\mathscr{P}r$ measure of the set of ω's for which $|(j_n(\omega)/n) - p| > \varepsilon$ approaches 0 as $n \to \infty$ for any $\varepsilon > 0$.

SLLN The $\mathscr{P}r$ measure of the set of all ω's such that $\lim_{n\to\infty}(j_n(\omega)/n) \neq p$ is 0.

To put (SLLN) in its positive form, the Pr probability is one that the limiting relative frequency of heads converges to p.

As indicated in section 7, a form of the weak law of large numbers can be formulated and proved without the help of countable additivity. Roughly, for any $\varepsilon > 0$, the probability (in the objective sense or in the degree-of-belief sense tempered by Lewis's principal principle) that the actually observed relative frequency of heads differs from p by more than ε goes to 0 as the number of flips goes to infinity. This form of the law of large numbers is to be found in the work of Bernoulli. The strong form of the law of large numbers, which requires countable additivity, was not proved until this century (see Billingsley 1979 for a proof).

3 Success Stories

The successes of the Bayesian approach to confirmation fall into two categories. First, there are the successes of Bayesianism in illuminating the virtues and pitfalls of various approaches to confirmation theory by providing a Bayesian rationale for what are regarded as sound methodological procedures and by revealing the infirmities of what are acknowledged as unsound procedures. The present chapter reviews some of these explanatory successes. Second, there are the successes in meeting a number of objections that have been hurled against Bayesianism. The following chapter discusses several of these successful defenses. Taken together, the combined success stories help to explain why many Bayesians display the confident complacency of true believers. Chapters 5 to 9 will challenge this complacency. But before turning to the challenges, let us give Bayesianism its due.

1 Qualitative Confirmation: The Hypotheticodeductive Method

When Carl Hempel published his seminal "Studies in the Logic of Confirmation" (1945), he saw his essay as a contribution to the logical empiricists' program of creating an inductive logic that would parallel and complement deductive logic. The program, he thought, was best carried out in three stages: the first stage would provide an explication of the qualitative concept of confirmation (as in 'E confirms H'); the second stage would tackle the comparative concept (as in 'E confirms H more than E' confirms H''); and the final stage would concern the quantitative concept (as in 'E confirms H to degree r'). In hindsight it seems clear (at least to Bayesians) that it is best to proceed the other way around: start with the quantitative concept and use it to analyze the comparative and qualitative notions. The difficulties inherent in Hempel's own account of qualitative confirmation will be studied in section 2. This section will be devoted to the more venerable hypotheticodeductive (HD) method.

The basic idea of HD methodology is deceptively simple. From the hypothesis H at issue and accepted background knowledge K, one deduces a consequence E that can be checked by observation or experiment. If Nature affirms that E is indeed the case, then H is said to be HD-confirmed, while if Nature affirms $\neg E$, H is said to be HD-disconfirmed. The critics of HD have so battered this account of theory testing that it would be unseemly to administer any further whipping to what is very nearly a dead horse.[1] Rather, I will review the results of the jolly Bayesian postmortem.

Suppose that (a) $\{H, K\} \models E$, (b) $0 < \Pr(H/K) < 1$, and (c) $0 < \Pr(E/K) < 1$.[2] Condition (a) is just the basic HD requirement for confirmation. Condition (b) says that on the basis of background knowledge K, H is not known to be almost surely true or to be almost surely false, and (c) says likewise for E. By Bayes's theorem and (a), it follows that

$$\Pr(H/E \ \& \ K) = \Pr(H/K)/\Pr(E/K). \tag{3.1}$$

By applying (b) and (c) to (3.1), we can conclude that $\Pr(H/E \ \& \ K) > \Pr(H/K)$, i.e., E *incrementally confirms* H relative to K. Thus Bayesianism is able to winnow a valid kernel of the HD method from its chaff.

(To digress, this alleged success story might be questioned on the grounds that HD testing typically satisfies not condition (a) but rather a condition Hempel calls the "prediction criterion" of confirmation; namely, (a') E is logically equivalent to $E_1 \ \& \ E_2$, $\{H, K, E_1\} \models E_2$, but $\{H, K\} \not\models E_2$. That is, HD condition (a) is satisfied with respect to the conditional prediction $E_1 \rightarrow E_2$, but the total evidence consists of E_1 and E_2 together. Let us use Bayes's theorem to draw out the consequences of (a'). It follows that $\Pr(H/E_1 \ \& \ E_2 \ \& \ K) = \Pr(H/E_1 \ \& \ K)/\Pr(E_2/E_1 \ \& \ K)$. Thus if $\Pr(E_2/E_1 \ \& \ K) < 1$ and $\Pr(H/E_1 \ \& \ K) = \Pr(H/K)$, the total evidence $E_1 \ \& \ E_2$ incrementally confirms H. These latter two conditions are satisfied in typical cases of HD testing. For example, let H be Newton's theory of planetary motion, let E_1 be the statement that a telescope is pointed in such and such a direction tomorrow at 3:00 P.M., and let E_2 be the statement that Mars will be seen through the telescope. Presumably, E_1 is probabilistically irrelevant to the theory, and E_2 is uncertain on the basis of E_1 and K.)

Notice also that from (3.1) it follows that the smaller the value of the prior likelihood $\Pr(E/K)$, the greater the incremental difference $\Pr(H/E \ \& \ K) - \Pr(E/K)$, which seems to validate the saying that the more surprising the evidence is, the more confirmational value it has. This observation, however, is double-edged, as we will see in chapter 5.

The problem of irrelevant conjunction, one of the main irritants of the HD method, is also illuminated. If $\{H, K\} \models E$, then also $\{H \ \& \ I, K\} \models E$, where I is anything you like, including a statement to which E is, intuitively speaking, irrelevant. But according to the HD account, E confirms $H \ \& \ I$. In a sense, the Bayesian analysis concurs, since if $\Pr(H \ \& \ I/K) > 0$, it follows from the reasoning above that E incrementally confirms $H \ \& \ I$. However, note that it follows from (3.1) that the amounts of incremental confirmation that H and $H \ \& \ I$ receive are proportional to their prior probabilities:

$$\Pr(H/E \ \& \ K) - \Pr(H/K) = \Pr(H/K)[(1/\Pr(E/K)) - 1]$$

$$\Pr(H \ \& \ I/E \ \& \ K) - \Pr(H \ \& \ I/K) = \Pr(H \ \& \ I/K)[(1/\Pr(E/K)) - 1].$$

Since in general $\Pr(H \ \& \ I/K) < \Pr(H/K)$, adding the irrelevant conjunct I to H lowers the incremental confirmation afforded by E.

Finally, it is worth considering in a bit more detail the case of HD disconfirmation. Thus, suppose that when Nature speaks, she pronounces $\neg E$. If $\{H, K\} \models E$ and if K is held to be knowledge, then H must be false, so HD disconfirmation would seem to be equivalent to falsification. But as Duhem and Quine have reminded us, the deduction of observationally decidable consequences from high-level scientific hypotheses often requires the help of one or more auxiliary assumptions A. It is not fair to ignore this problem by sweeping the A's under the rug of K, since the A's are often every bit as questionable as H itself. Thus from Nature's pronouncement of $\neg E$ all that can be concluded from deductive logic alone is that $\neg H \vee \neg A$. If HD methodology were all there is to inductive reasoning, then there would be no principled way to parcel out the blame for the false prediction, and we would be well on the way to Duhem and Quine holism (see section 4 below). In particular, H could be maintained come what may if the only constraints operating were those that followed from direct observation and deductive logic. But the fact that the majority of scientists sometimes regard the maintenance of a hypothesis as reasonable and sometimes not is a fact of actual scientific practice that cries out for explanation. The Bayesian attempt at an explanation will be examined in section 7 below.

2 Hempel's Instance Confirmation

Having rejected the HD or prediction criterion of confirmation, Hempel constructed his own analysis of qualitative confirmation on a very different basis. He started with a number of conditions that he felt that any adequate theory of confirmation should satisfy, among which are the following:

Consequence condition If $E \models H$, then E confirms H.

Consistency condition If E confirms H and also H', then $\not\models \neg(H \& H')$.

Special consequence condition If E confirms H and $H \models H'$, then E confirms H'.

Hempel specifically rejected the converse consequence condition:

Converse consequence condition If E confirms H and $H' \models H$, then E confirms H'.

For to add the last condition to the first three would lead to the disaster that any E confirms any H.[3] (Note that HD confirmation satisfies the converse consequence condition but violates both the consistency condition and the special consequence condition.)

Hempel's basic idea for finding a definition of qualitative confirmation satisfying his adequacy conditions was that a hypothesis is confirmed by its positive instances. This seemingly simple and straightforward notion turns out to be notoriously difficult to pin down.[4] Hempel's own explication utilized the notion of the *development* of a hypothesis for a finite set I of individuals. Intuitively, $\text{dev}_I(H)$ is what H asserts about a domain consisting of just the individuals in I. Formally, $\text{dev}_I(H)$ for a quantified H is arrived at by peeling off universal quantifiers in favor of conjunctions over I and existential quantifiers in favor of disjunctions over I. Thus, for example, if $I = \{a, b\}$ and H is $(\forall x)(\exists y)Lxy$ (e.g., "Everybody loves somebody"), $\text{dev}_I(H)$ is $(Laa \vee Lab) \& (Lbb \vee Lba)$. We are now in a position to state the main definitions that constitute Hempel's account.

Definition *E directly Hempel-confirms H* iff $E \models \text{dev}_I(H)$, where I is the class of individuals mentioned in E.

Definition *E Hempel-confirms H* iff there is a class C of sentences such that $C \models H$ and E directly confirms each member of C.[5]

Definition *E Hempel-disconfirms H* iff E Hempel-confirms $\neg H$.

The difficulties with Hempel's account can be grouped into three categories. The first concerns the pillars on which the account was built: Hempel's so-called adequacy conditions. Bayesians have at least two ways of defining qualitative confirmation, one of which we already encountered in section 1; namely, *E incrementally confirms H* relative to K iff $\Pr(H/E \& K) > \Pr(H/K)$. The second is an absolute rather than incremental notion; specifically, *E absolutely confirms H* relative to K iff $\Pr(H/E \& K) \geqslant k > .5$. (A third criterion sometimes used in the literature, e.g., Mackie 1963, says that E confirms H relative to K just in case $\Pr(E/H \& K) > \Pr(E/K)$. The reader can easily show that on the assumption that none of the probabilities involved is zero, this *likelihood criterion* is equivalent to the incremental criterion.) In both instances there appears to be a mismatch, since Hempel's account is concerned with a two-place relation 'E confirms H' rather than with a three-place relation ('E confirms H relative to K'). The Bayesians can accommodate themselves to Hempel either by taking K to be empty or by supposing that K has been learned and then working with the new probability function $\Pr'(\cdot) = \Pr(\cdot/K)$ obtained by conditionalization. But since one of the morals the Bayesians want to draw is that background knowledge can make a crucial difference to confirmation, I will continue to make K an explicit factor in the confirmation equation.

The first difficulty for Hempel's account can now be stated as a dilemma. For any choice of K compatible with H, Hempel's adequacy conditions accord well with the absolute notion of Bayesian confirmation. For example, if $\Pr(H/E \& K) > .5$ and $H \models H'$, then $\Pr(H'/E \& K) > .5$, so the special consequence condition is satisfied. But absolute confirmation cannot be what Hempel had in mind, since he holds that the observation of a single black raven a confirms the hypothesis that all ravens are black, even though for typical K's, $\Pr((\forall x)(Rx \rightarrow Bx)/Ra \& Ba \& K) \ll .5$. On the other hand, while the incremental concept of confirmation allows that a single instance can confirm a general hypothesis, both the consistency condition and the special consequence condition fail for not atypical K's, as examples by Carnap (1950) and Salmon (1975) show.[6] Of course, there may be some third probabilistic condition of confirmation that allows Hempel's account to pass between the horns of this dilemma. But it is up to the defender of Hempel's instance confirmation to produce the *tertium quid*. And even to conduct the search for a probabilistic *tertium quid* is to fall into the hands of the Bayesians.

The second category of difficulties revolves around the question of whether Hempel's account is too narrow. One reason for thinking so is that, as Hempel himself notes, a hypothesis of the form

$$(\forall x)(\exists y)Rxy \& (\forall x)(\forall y)(\forall z)[(Rxy \& Ryz) \rightarrow Rxz] \& (\forall x)\neg Rxx$$

cannot be Hempel-confirmed by any consistent E, since the development of such a hypothesis for a finite domain is inconsistent. Nor is the hypothesis $(\forall x)(\forall y)Rxy$ Hempel-confirmed by the set of evidence statements $\{Ra_ia_j\}$, where $i = 1, 2, \ldots, 10^9$ and $j = 1, 2, \ldots, 10^9 - 1$. Even more troublesome is the fact that Hempel's account is silent about how theoretical hypotheses are confirmed, for if, as Hempel intended, E is stated purely in the observational vocabulary and if H is stated in a theoretical vocabulary disjoint from the observational vocabulary, then E cannot, except in very uninteresting cases, Hempel-confirm H.[7] This silence is a high price to pay for overcoming some of the defects of the more vocal HD method.

Clark Glymour (1980) has sought to preserve Hempel's idea that hypotheses are confirmed by deducing positive instances of them from observation reports. In the case where H is stated in theoretical vocabulary, Glymour's bootstrapping method allows the deduction to proceed via auxiliary hypotheses, typically drawn from a theory T of which H itself is a part.[8] His basic confirmation relation is thus three-place: E confirms H relative to T.

The Bayesian response to these difficulties and to Glymour's reaction to them is twofold. First, there is no insuperable problem about how observational data can confirm, in either the incremental or absolute sense, a theoretical hypothesis; indeed, the application of Bayes's theorem shows just how such confirmation takes place, at least on the assumption that the prior probability of the hypothesis is nonzero (a matter that will be taken up in chapter 4). Second, unless bootstrap confirmation connects to reasons for believing the hypothesis or theory, it is of no interest. But once the connection is made, the bootstraps can be ignored in favor of the standard Bayesian account of reasons to believe. This matter will be examined in more detail in section 4 below.

The third category of difficulties is orthogonal to the second. Now the worry is that while Hempel's instance confirmation may be too narrow in some respects, it may be too liberal in other respects. Consider again the ravens hypothesis: $(\forall x)(Rx \rightarrow Bx)$. Which of the following evidence statements Hempel-confirm it?

E_1: $Ra_1 \,\&\, Ba_1$

E_2: $\neg Ra_2 \,\&\, \neg Ba_2$

E_3: $\neg Ra_3$

E_4: Ba_4

E_5: $\neg Ra_5 \,\&\, Ba_5$

E_6: $Ra_6 \,\&\, \neg Ba_6$

Only E_6 fails to Hempel-confirm the hypothesis, and that is because E_6 falsifies it. The indoor ornithology involved in using E_2 to E_5 as confirmation of the ravens hypothesis has struck many commentators as too easy to be correct. Bayesian treatments of Hempel's ravens paradox will be taken up in the following section.

If anything is safe in this area, it would seem to be that E_1 does confirm $(\forall x)(Rx \rightarrow Bx)$. But safe is not sure. Recall that Hempel's definition of confirmation is purely syntactical in that it is neutral to the intended interpretation of the predicates. This means that E_1 Hempel-confirms $(\forall x)(Rx \rightarrow Bx)$ even if we take Bx to mean not that x is black but that x is blite, i.e., x is first examined before the year 2000 and is black, or else is not examined before 2000 and is white. Let a_i be first examined in the year i. Then by the special consequence condition, $Ra_1 \,\&\, Ba_1 \,\&\, Ra_2 \,\&\, Ba_2 \,\&\, \ldots \,\&\, Ra_{1999} \,\&\, Ba_{1999}$ Hempel-confirms the prediction $Ra_{2001} \rightarrow Ba_{2001}$, i.e., the prediction that if a_{2001} is a raven, then it is white, which is, to say the least, counterintuitive. We have here an instance of what Goodman (1983) calls the "new riddle of induction." The Bayesian treatment of this problem will be given in detail in chapter 4. But for now I will simply note on behalf of the Bayesians that they are not committed to assigning probabilities purely on the basis of the syntax of the hypothesis and the evidence, as Hempel's analogy between deductive and inductive logic would suggest. The present example is enough to show that an adequate account of confirmation must be sensitive to semantics, and this lesson is easily incorporated into Bayesianism.

3 The Ravens Paradox

In sections 1 and 2 Bayesianism gained reflected glory of sorts from the whippings the HD and Hempel accounts took. It is time for Bayesianism to earn additional glory of a more positive sort.

Hempel took it as a desirable consequence of his account that the evidence $Ra \,\&\, Ba$ confirms the hypothesis $(\forall x)(Rx \rightarrow Bx)$.[9] The paradox of the ravens in one of its forms arises from the fact that on Hempel's analysis, the evidence $\neg Rb \,\&\, \neg Bb$ also confirms $(\forall x)(Rx \rightarrow Bx)$. Before turning to the Bayesian analysis of the paradox itself, it is worth noting that the Bayesian is not even willing to go the first step with Hempel without first looking both ways.

Suppose that $0 < \Pr(H/K) < 1$, where H stands for the ravens hypothesis. Then by an application of Bayes's theorem it follows that finding a to be a black raven induces incremental confirmation,

$$\Pr(H/Ra \,\&\, Ba \,\&\, K) > \Pr(H/K),$$

just in case

$$\Pr(Ra/H \,\&\, K) > \Pr(Ra/\neg H \,\&\, K) \times \Pr(Ba/Ra \,\&\, \neg H \,\&\, K).$$

Incremental disconfirmation results just in case the inequality is reversed.[10] The reader is invited to reflect on the kinds of background knowledge K that will make or break these inequalities. Consider, for instance, a version of I. J. Good's (1967) example. We are supposed to know in advance (K) that we belong to one of two bird universes: one where there are 100 black ravens, no nonblack ravens, and 1 million other birds, or else one where there are 1,000 black ravens, 1 white raven, and 1 million other birds. Bird a is selected at random from all the birds and found to be a black raven. This evidence, Good claims, undermines the ravens hypothesis. Use the above formula to test this claim. Such exercises help to drive home the point that a two-place confirmation relation that ignores background evidence is not very useful.

Let us turn now to the Bayesian treatment of the bearing of the evidence of nonblack nonravens on the ravens hypothesis. Suppes (1966) invites us to consider an object a drawn at random from the universe. Set

$$\Pr(Ra \,\&\, Ba/K) = p_1, \qquad \Pr(Ra \,\&\, \neg Ba/K) = p_2,$$
$$\Pr(\neg Ra \,\&\, Ba/K) = p_3, \qquad \Pr(\neg Ra \,\&\, \neg Ba/K) = p_4. \tag{3.2}$$

Then

$$\Pr(\neg Ba/Ra \,\&\, K) = p_2/(p_1 + p_2) \tag{3.3}$$

and

$$\Pr(Ra/\neg Ba \,\&\, K) = p_2/(p_2 + p_4). \tag{3.4}$$

From (3.3) and (3.4) it follows that $\Pr(\neg Ba/Ra \,\&\, K) > \Pr(Ra/\neg Ba \,\&\, K)$ iff $p_4 > p_1$. But from what we know of the makeup of our universe, it seems safe to assume that $p_4 \gg p_1$, with the consequence that the conditional probability of a's being nonblack, given that it is a raven, is much greater than the conditional probability of a's being a raven, given that it is nonblack. The moral Suppes wants us to draw from this is that sampling from the class of ravens is more productive than sampling from the class of nonblack objects, since the former procedure is more likely to produce a counterexample to the ravens hypothesis.

There are two qualms about this moral. The first is that it doesn't seem directly useful to Bayesians; indeed, at first blush it seems more congenial to a Popperian line that emphasizes the virtues of attempted falsifications of hypotheses. Second, it is not clear how the moral follows from the inequality derived, since a was supposed to result from a random sample of the universe at large rather than from a random sample of either the class of ravens or the class of nonblack objects.

Horwich's (1982) attack on the ravens paradox starts from the observation that there are several ways to obtain the evidence $Ra \,\&\, Ba$, namely, to pick an object at random from the universe at large and find that it has both ravenhood and blackness, to pick an object at random from the class of ravens and find that it is black, or to pick an object at random from the class of black things and find that it is a raven. A similar remark applies to the evidence $\neg Rb \,\&\, \neg Bb$. Horwich introduces the notation R^*a to mean that a was drawn at random from the class of ravens and the notation $\neg B^*b$ to mean that b was drawn at random from the class of nonblack things. To illuminate the ravens paradox, he wants to compare the confirmational effects of the two pieces of evidence $R^*a \,\&\, Ba$ and $\neg B^*b \,\&\, \neg Rb$. According to Horwich's application of Bayes's theorem,

$$\Pr(H/R^*a \,\&\, Ba \,\&\, K) = \Pr(H/K)/\Pr(R^*a \,\&\, Ba/K) \tag{3.5}$$

and

$$\Pr(H/\neg B^*b \,\&\, \neg Rb \,\&\, K) = \Pr(H/K)/\Pr(\neg B^*b \,\&\, \neg Rb/K), \tag{3.6}$$

where K is the same as before. Thus

$$\Pr(H/R^*a \,\&\, Ba \,\&\, K) > \Pr(H/\neg B^*b \,\&\, \neg Rb \,\&\, K)$$
$$\text{iff} \quad \Pr(\neg B^*b \,\&\, \neg Rb/K) > \Pr(R^*a \,\&\, Ba/K).$$

But the latter is true for our universe, Horwich asserts.

But as with Suppes's construction, it is not clear how this conclusion follows. In the first place, why is it true (as (3.5) and (3.6) assume) that

$$\Pr(R^*a \,\&\, Ba/H \,\&\, K) = \Pr(\neg B^*b \,\&\, \neg Rb/H \,\&\, K) = 1?$$

It is true that the probability of a randomly chosen raven being black, given H & K, is 1, but $\Pr(R^*a \,\&\, Ba/H \,\&\, K)$ is the probability that an object a is randomly chosen from the class of ravens and is black, given H & K, and this probability is surely not 1. In the second place, comparing $\Pr(\neg B^*b \,\&\, \neg Rb/K)$ and $\Pr(R^*a \,\&\, Ba/K)$ involves a comparison of the probability that an object will be randomly sampled from the class of ravens with the probability that it will be randomly sampled from the class of nonblack things, and such a comparison seems peripheral to the paradox at best.

Horwich's basic idea can be brought to fruition by putting into the background knowledge $\hat{K}$ the information that R^*a and $\neg B^*b$. Bayes's theorem can then be legitimately applied to the new $\hat{K}$ to conclude that

$$\Pr(H/Ra \,\&\, Ba \,\&\, \hat{K}) = \Pr(H/\hat{K})/\Pr(Ba/\hat{K}) \tag{3.7}$$

and

$$\Pr(H/\neg Rb \,\&\, \neg Rb \,\&\, \hat{K}) = \Pr(H/\hat{K})/\Pr(\neg Rb/\hat{K}). \tag{3.8}$$

Thus, relative to this $\hat{K}$, the evidence Ra & Ba has more confirmational value vis-à-vis the ravens hypothesis than does $\neg Rb$ & $\neg Bb$ just in case $\Pr(\neg Rb/\hat{K}) > \Pr(Ba/\hat{K})$. A further application of the principle of total probability shows that this latter inequality holds just in case $\Pr(\neg Ba/\neg H \,\&\, \hat{K}) > \Pr(Rb/\neg H \,\&\, \hat{K})$. This last inequality presumably does hold in our universe, for given that some ravens are nonblack ($\neg H$), we are more likely to produce one of them by sampling from the class of ravens than by sampling from the class of nonblack things simply because of the known size and heterogeneity of the class of nonblack things as compared with the known size of the class of ravens. Suppes is thus vindicated after all, since the greater confirmatory power of Ra & Ba over $\neg Rb$ & $\neg Bb$ has to do with the relative threats of falsification. In this way Bayesianism pays a backhanded compliment to Popper's methodology; namely, it is precisely because, contrary to Popper, inductivism is possible that the virtues of sincere attempts to falsify can be recognized.[11]

Similar points are made by Gaifman (1979), although his assumed sampling procedure is somewhat different. Let $\tilde{K}$ report that c was drawn at random from the universe and found to be a raven and that d was also drawn at random from the universe and found to be nonblack. An analysis like the one above shows that

$$\Pr(H/Rc \,\&\, Bc \,\&\, \tilde{K}) > \Pr(H/\neg Rd \,\&\, \neg Bd \,\&\, \tilde{K})$$

just in case $\Pr(\neg Bc/\neg H \,\&\, \tilde{K}) > \Pr(Rd/\neg H \,\&\, \tilde{K})$.

But the procedure of sampling from the universe at large can be wasteful, since it can produce relatively useless results, such as $\neg Re$ & Be. Moreover, one can wonder whether the evidence Ra & Ba, under the assumption that a was drawn at random from the class of ravens, gives better confirmational value than the evidence Rc & Bc, under the assumption that c was drawn at random from the universe at large, i.e., whether

$$\Pr(H/Ra \,\&\, Ba \,\&\, \hat{K} \,\&\, \tilde{K}) > \Pr(H/Rc \,\&\, Bc \,\&\, \hat{K} \,\&\, \tilde{K}).$$

I leave it to the reader to ponder this question with the clue that the answer is positive just in case

$$\Pr(\neg Ba/\neg H \,\&\, \hat{K} \,\&\, \tilde{K}) > \Pr(\neg Bc/\neg H \,\&\, \hat{K} \,\&\, \tilde{K}).$$ [12]

4 Bootstrapping and Relevance Relations

In *Theory and Evidence* (1980) Glymour saw bootstrapping relations not only as a means of extending Hempel's instance confirmation to theoretical hypotheses but also as an antidote to Duhem and Quine holism. It makes a nice sound when it rolls off the tongue to say that our claims about the physical world face the tribunal of experience not individually but only as a corporate body. But scientists, no less than business executives, do not typically act as if they are at a loss as to how to distribute praise through the corporate body when the tribunal says yea, or blame when the tribunal says nay. This is not to say that there is always a single correct way to make the distribution, but it is to say that in many cases there are firm intuitions. Bootstrap relations would help to explain these intuitions if they helped to explain why it is that for some but not all H's that are part of a theory T, E bootstrap-confirms H relative to T.

As a sometime Bayesian I now think that bootstrapping should be abandoned in favor of a Bayesian analysis. Bayesians can be sympathetic to the two motivations for bootstrapping mentioned above in section 2. At