Bayesian Confirmation Theory: Inductive Logic, or Mere Inductive Framework?

Michael Strevens

Synthese 141:365–379, 2004

Abstract

Does the Bayesian theory of confirmation put real constraints on our inductive behavior? Or is it just a framework for systematizing whatever kind of inductive behavior we prefer? Colin Howson (*Hume's Problem*) has recently championed the second view. I argue that he is wrong, in that the Bayesian apparatus as it is usually deployed does constrain our judgments of inductive import, but also that he is right, in that the source of Bayesianism's inductive prescriptions is not the Bayesian machinery itself, but rather what David Lewis calls the "Principal Principle".

1. Inductive Logics versus Inductive Frameworks

In *Hume's Problem*, Colin Howson asks whether Bayesian confirmation theory (BCT) solves the problem of induction (Howson 2001). His answer is that it does not. But Howson is, of course, an avowed Bayesian, and he wants to justify BCT all the same. What BCT offers, Howson claims, is not an inductive logic so much as an inductive framework; as a framework, he continues, BCT concerns only matters of internal consistency and so can be justified. Of these two claims, the present paper will be concerned with the first, that BCT is a mere framework, not a logic; questions of justification will be put to one side.

What is the difference between an inductive logic and an inductive framework? (These are my terms, not Howson's.)[1] Inherent in an inductive logic are certain inductive commitments, for example, a commitment to the proposition that the future resembles the past. An inductive logic tells us how to reason in accordance with these commitments. I use the term *inductive logic* in the broadest possible sense, then, so as to include any system for ampliative inference.

An inductive framework, by contrast, has no intrinsic inductive commitments. The user of a framework supplies their own commitments; what the framework does is to provide an apparatus for transforming any given set of inductive commitments into a full-fledged inductive reasoning system. Once you incorporate your favorite inductive commitments into an inductive framework, then, you get an inductive logic. What is the framework doing? The purpose of an inductive framework, according to Howson, is to ensure that you apply your inductive commitments consistently to every piece of evidence.

---

1. Howson proposes that the term *inductive logic* be used to refer to the a priori component of any inductive system. Because, on his view, no inductive commitments can be known a priori, it will turn out that an inductive logic just is an inductive framework.

Howson's distinction brings to mind the inductive framework developed by Carnap in *Logical Foundations of Probability* (Carnap 1950). Carnap's framework has a single parameter λ that can range from zero to infinity. Different values of λ correspond to different inductive commitments. Setting λ to zero yields the straight rule of induction (where a probability for an event type is set equal to its observed frequency). Making λ infinitely large yields a policy on which evidence is ignored (Carnap's $c\dagger$). Setting λ equal to 2 yields Carnap's c^*, equivalent to Laplace's rule of succession.

The logic/framework distinction constitutes more of a spectrum than a dichotomy. At one end of the spectrum is a pure inductive framework, making no inductive commitments at all. As you add stronger and stronger inductive commitments, you move along the spectrum, until at the other end you have a system that choreographs precisely your every inductive move. In practice, the extremes are rare. Inductive logics tend to allow some sort of freedom in setting up the system; Carnap's inductive logics, for example, depend very much on a choice of language. Similarly, inductive frameworks tend not to be compatible with just any inductive commitment; they thereby incorporate a certain low level of inductive commitment themselves. Again, Carnap's system provides an example. (For the inductive commitments of even very weak, framework-like Bayesianisms, see the end of section 4.)

Hume's Problem is built on the claim that BCT is an inductive framework. Yet it is standard to interpret BCT as being a kind of inductive logic, not a mere framework. I pose two questions in this paper. First, is Howson correct that BCT is a mere framework, or does it house stronger inductive commitments than he supposes? Second, if Howson is wrong, as I will argue, what are the sources of BCT's commitments?

Before I continue, let me say something more about the nature of inductive commitments. An inductive commitment is either a grand rule stating which hypotheses ought to be preferred, given certain kinds of evidence, or a grand generalization about the nature of the universe that provides the foundation for such a rule.

I count the following, for example, as inductive commitments:

1. Favor hypotheses that predict more of the same (a grand rule).
2. The principle of the uniformity of nature (a grand generalization),
3. Favor hypotheses that entail the observed evidence.
4. Favor those hypotheses with the higher physical likelihoods, that is, those hypotheses that assign relatively higher physical probabilities to the evidence. This is the *likelihood lover's principle*. (It is not to be confused with what philosophers of statistics call the likelihood principle, a much stronger and more controversial principle.)
5. Favor hypotheses that provide better explanations of the evidence.
6. Favor hypotheses phrased in terms of predicates like *green* over theories phrased in terms of predicates like *grue*, all other things being equal.
7. The universe is governed by relatively simple principles (and so you should favor simple hypotheses over complex hypotheses, all other things being equal).
8. The universe is governed by beautiful principles (and so you should favor beautiful hypotheses over ugly hypotheses, all other things being equal).

When Howson says that BCT is an inductive framework, he means that it incorporates few or no inductive commitments. What are his reasons for holding this view? His argument has the following general form.

1. The inductions recommended by BCT depend in part on certain of the scientist's subjective probabilities, called the *priors*,
2. The priors are constrained only very weakly, and
3. Depending on how the priors are set within these very weak constraints, BCT will implement any number of different, competing in-

ductive commitments.

In short, the constraints on subjective probabilities are not strong enough to limit BCT to any one set of inductive commitments. Or—to put the point more positively—you can incorporate almost any inductive commitments you like into BCT just by choosing your priors appropriately.

Of the premises of Howson's argument, (1) is uncontroversial, and (2), though disputed by some Bayesians (see the end of section 2), is allowed by many, certainly the majority of contemporary Bayesians. I want to ask whether Howson is right in asserting (3).

Howson's argument for (3) in *Hume's Problem* seems to lie for the most part in chapter four, which explains the inability of BCT or any similar probabilistic method to solve the grue problem. The difficulties created by *grue* lead Howson to the conclusion that, unless some discrimination against "grueish" hypotheses is made in the priors, the observed evidence can never, in virtue of BCT alone, warrant a particular expectation about the future. (The form of the argument is sketched at the end of section 3 of this paper, once the Bayesian apparatus has been introduced.) A similar argument has been made by Albert (2001).[2]

It does not follow, however, that BCT entirely lacks inductive commitment. Even assuming that Howson's or Albert's arguments can be generalized to evidence of all varieties, it may be that

1. There is an inductive commitment inherent in BCT that is independent of the grue problem, and that is evident in Bayesian reasoning whether or not the grue problem is resolved. Or it may be that
2. There is a latent inductive commitment in BCT that is not evident as long as the grue problem is left open, but that shows itself once you

2. Albert's conclusion is rather stronger than Howson's; he infers that BCT constrains us not at all. Howson, noting that it is quite possible to violate Bayes' rule, argues that there are constraints, but that these do not amount to an inductive commitment.

have taken some stance on *grue* by putting a prior probability distribution over the possible hypotheses, both grueish and non-grueish (presumably favoring, in most cases, the non-grueish). The commitment would not, of course, be introduced by your prior probability distribution—that would be a vindication of Howson and Albert. Rather, your priors would merely enable the implementation of a pre-existing commitment, perhaps simply by giving the Bayesian apparatus something to work with.

What Howson and Albert have reason to conclude, then, is that the inductive commitments of BCT are not wholly sufficient in themselves to licence particular predictions given particular sets of evidence. That is a significant conclusion—perhaps it is all the conclusion that Howson really wants—but it does not entail that BCT harbors no inductive commitments whatsoever. I aim to continue the search for inductive commitments in BCT, in particular, the search for those commitments consistent with Howson's and Albert's conclusions.

I will limit my discussion to one particular version of BCT, which I call modern Bayesianism. Most contemporary proponents of BCT subscribe to some form of modern Bayesianism. Recent influential presentations of modern Bayesianism can be found in John Earman's *Bayes or Bust?* and Howson and Urbach's *Scientific Reasoning: The Bayesian Approach* (Earman 1992; Howson and Urbach 1993). Earman's version is set out in a chapter titled "The Machinery of Modern Bayesianism", from which I have taken the name. (I should note that the outlines of modern Bayesianism can be discerned, according to Earman, even in Thomas Bayes' original paper; it is not, then, *merely* modern.) For comments on alternatives to modern Bayesianism, see the end of section 2.

Modern Bayesianism does, I will argue, incorporate and implement serious inductive commitments. In particular, it implements the likelihood lover's principle. But on closer inspection, it turns out that this commitment is not inherent in the part of modern Bayesianism that is properly Bayesian, but in a separate component of modern Bayesianism that I call the probability coordination principle (my generic name for rules such as David Lewis's Principal Principle). This is unexpected, because the probability coordination principle has not, traditionally, been thought of as embodying any kind of inductive commitment. Howson himself, to take an especially telling example, while holding that inductive commitments cannot be given an a priori justification, has attempted to provide an a priori justification for the probability coordination principle (Howson and Urbach 1993, 344–345).[3]

2. Modern Bayesianism

At the core of modern Bayesianism is a rule for changing the subjective probabilities assigned to hypotheses in the light of new evidence. This rule is Bayes' rule, which states that, on encountering some piece of evidence e, you should change your subjective probability for each hypothesis h to your old probability for h conditional on e. In symbols,

$$C^{+}(h) = C(h|e),$$

where $C(\cdot)$ is your subjective probability distribution before observing e and $C^{+}(\cdot)$ is your subjective probability distribution after observing e.

It follows from the definition of conditional probability that

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

3. In *Hume's Problem*, Howson's attitude towards the epistemic status of the probability coordination principle is more guarded (pp. 237–8). Given the brevity of his comments, it is hard to say whether or not he has abandoned his earlier view.

This result, Bayes' theorem, can be used to give the following far more suggestive formulation of Bayes' rule:

$$C^{+}(h) = \frac{C(e|h)}{C(e)} C(h).$$

More or less anyone who counts themselves a proponent of BCT thinks that this rule is *the* rule that governs the way that scientists' opinions should change in the light of new evidence.

All the terms in the rule are, on the Bayesian interpretation, subjective probabilities, reflecting psychological facts about the scientist rather than observer-independent truths about h and e. This has led to the accusation that BCT is far too subjective to be a serious contender as an account of the confirmation of scientific theories.

Modern Bayesianism attempts to reply to this accusation not by eliminating the subjectivity of Bayesian conditionalization altogether, but by concentrating the subjectivity in just one set of subjective probabilities, namely, the subjective probabilities that a scientist has for the different hypotheses before any evidence comes in. These are all probabilities of the form $C(h)$. They are what are called the *prior probabilities*, or the *priors* for short. What modern Bayesianism sets out to do, then, is to show that there are objective constraints on the assignment of the other subjective probabilities in the conditionalization rule, namely, the probabilities of the form $C(e|h)$, called the subjective likelihoods, and $C(e)$.

Modern Bayesianism puts an objective constraint on the subjective likelihoods by requiring that the subjective probability of some piece of evidence e given some hypothesis h be set equal to the physical probability that h ascribes to e. Modern Bayesianism is committed, then, to the following rule, sometimes called Miller's Principle after David Miller (1966):

$$C(e|h) = P_h(e),$$

where $P_h(\cdot)$ is the physical probability distribution posited by the hypothesis h. Miller's Principle is a particularly simple version of the rule; a more sophisticated version is worked out in Lewis (1980). Lewis called his rule the Principal Principle; he later decided that the Principal Principle is incorrect and endorsed what he called the New Principle (Lewis 1994).[4] It is useful to have a generic name for principles of this sort. I call them *probability coordination principles*. What modern Bayesianism uses to objectify the subjective likelihoods $C(e|h)$, then, is some probability coordination principle or other. It does not matter, for my purposes, which particular principle is fixed upon; I will refer to whichever one is chosen as *the* probability coordination principle, or PCP.

Note that this objective fixing of the subjective likelihoods assumes that all competing hypotheses are probabilistic theories that range over events of the types instantiated by the evidence, so that each competing hypothesis assigns a definite physical probability to any possible piece of evidence. For the sake of the argument, I will grant this assumption.

Modern Bayesianism puts an objective constraint on the subjective probability $C(e)$ for the evidence by way of a theorem of the probability calculus, the theorem of total probability, which asserts in one variant that:

$$C(e) = C(e|h_1)C(h_1) + \cdots + C(e|h_n)C(h_n),$$

where the h_i are a complete set of competing hypotheses, assumed to be mutually exclusive and exhaustive. This gives a formula for $C(e)$ in terms of the subjective likelihoods $C(e|h_i)$, which are objectively constrained by PCP,

4. The issue on which the principles differ is the handling of certain kinds of (fairly esoteric) information that defeat the application of Miller's simple principle. If you possess such *inadmissible information* (Lewis's term), you ought *not* to set your subjective probability for an event equal to the corresponding physical probability. The question is, first, what, if any, information counts as inadmissible, and second, how it should affect the relevant subjective probabilities. It is generally agreed that the problems arising from the existence of inadmissible information do not affect the day-to-day workings of BCT. For my views on this topic, see Strevens (1995).

and the priors $C(h_i)$. The priors are not constrained, but the total probability theorem reduces the subjectivity in $C(e)$ to the subjectivity in the priors, which modern Bayesianism already concedes. Given PCP and the theorem of total probability, then, the only subjective element in modern Bayesianism is the assignment of prior probabilities to the competing hypotheses: once you have chosen your priors, everything else is determined for you by PCP and the theorem of total probability.

In order to use the total probability theorem in this way, you must know the content of all the competing hypotheses. There cannot be an h_i in the application of the theorem that stands for the possibility of some unknown hypothesis being the correct one, because $C(e|h)$ for that hypothesis would not be fixed by PCP, leaving $C(e)$ under-constrained. This is an even stronger assumption than the assumption that all hypotheses ascribe a definite physical probability to e, but again, for the sake of the argument, I will not dispute it.

Because my focus in the following, most important parts of the paper is on modern Bayesianism exclusively, let me discuss briefly some other forms of BCT, so as to point out in passing their approximate location in a broader discussion of inductive commitment.

First, consider Bayesianisms weaker than modern Bayesianism. Modern Bayesianism without the probability coordination principle will be discussed in section 4, where I claim it is almost a pure framework. Any weaker Bayesianism will, if I am right, be at least as inductively uncommitted; an example would be BCT without Bayesian conditionalization itself, the effective result of a policy allowing you to reconsider your prior probabilities at any time (Levi 1980).

Second, consider Bayesianisms that strengthen modern Bayesianism by putting serious constraints on, and sometimes even uniquely determining, values for your prior probabilities. In so doing, these systems stand to make stronger inductive commitments than modern Bayesianism (though none that I know of clearly prescribes a resolution of the grue problem that would satisfy Howson and Albert).

Examples of strong Bayesianisms include those based on an a priori symmetry principle for the priors (Jaynes 1983); "empirical" Bayesianisms that use observed frequencies to calibrate priors, which in their adherence to calibration clearly make a certain inductive commitment (Dawid 1982); and "logical" Bayesianisms in which subjective likelihoods, above and beyond those that fall within the scope of PCP, are constrained by objective facts about "inductive support", also a clear inductive commitment (Keynes 1921). (You might think of these different Bayesianisms, weak and strong, as different ways of fleshing out an inductive framework even sparer than the one that Howson has in mind, but that is not my strategy here.)

3. Modern Bayesianism as Inductive Logic

Now I turn to the question whether modern Bayesianism makes any significant inductive commitments, that is, whether it is an inductive logic in its own right, or merely a framework for inductive logic, as Howson claims.

On the framework view, inductive commitments are added to modern Bayesianism by particular choices of priors. This seems rather odd: the priors look like opinions about particular hypotheses, not about the proper way to do induction.

Nevertheless, it is generally agreed that certain inductive commitments reside in the priors, and therefore that modern Bayesianism does not in itself either endorse or reject these commitments. With respect to certain commitments, then, the consensus is that modern Bayesianism acts like a framework. The commitments include the following:

1. Favor hypotheses phrased in terms of predicates like *green* over theories phrased in terms of predicates like *grue*, all other things being

equal. (As Howson and others have argued, modern Bayesianism does not discriminate against *grue*; any bias in favor of *green* over *grue* must be present in the priors, in the sense that hypotheses using *green* are assigned higher prior probabilities than their counterparts using *grue*.)

2. Favor simple hypotheses over complex hypotheses, all other things being equal. (If two hypotheses, one simple and one complex, assign the same probabilities to all observable phenomena, the apparatus of modern Bayesianism will not in itself favor one over the other. Any bias in favor of the simpler hypothesis must be present in the priors.)

Some inductive commitments are, however, inherent in the machinery of modern Bayesianism. To adopt the machinery is to commit yourself to the following inductive maxims:

1. Favor theories that entail the observed evidence, and
2. The likelihood lover's principle: favor theories that assign relatively higher physical probabilities to the evidence.

Of these two commitments, I want to focus on the second, the likelihood lover's principle, or LLP, which is sufficiently broad in its inductive scope, I submit, to place BCT towards the inductive logic end of the logic/framework spectrum.

To see that modern Bayesianism directs us to favor hypotheses with higher physical likelihoods, consider the Bayesian conditionalization rule with the subjective likelihood replaced by the corresponding physical probability, as required by PCP:[5]

$$C^{+}(h) = \frac{P_h(e)}{C(e)}C(h)$$

5. Assuming that there is no inadmissible information; see note 4.

Since $C(e)$ is the same for every hypothesis, each hypothesis h gets a probability boost that is proportional to the physical probability that it assigns to the evidence. Thus hypotheses with higher likelihoods will be relatively favored.

Modern Bayesianism takes advantage of this fact to show that, even if scientists disagree at the outset about the prospects of different hypotheses, their opinion will very likely eventually converge; the convergence is, of course, on the hypotheses that ascribe the highest probability to the evidence.

Convergence results—and applications of the likelihood lover's principle in general—cannot, however, be used to discriminate between hypotheses that assign the same probability to all the observed data. This is a key premise of Howson's and Albert's argument that Bayesianism has no inductive commitments.[6] The other key element is a method for constructing hypotheses that agree on all the evidence so far observed, but that disagree on the next piece of evidence. If you are to choose between these conflicting predictions, it can only be your prior probability distribution, and not Bayesian conditionalization, that inclines you one way or the other. Conditionalization alone does not recommend one prediction over the others.

But even if Howson and Albert are correct that the machinery of modern Bayesianism does not mandate particular predictions from particular data sets, it does not follow that modern Bayesianism enforces no inductive preferences whatsoever. The possibility envisaged at the end of section 1, that BCT has inductive commitments, but that the commitments do not, on their own, dictate definite predictions, turns out to be actual: the likelihood lover's principle is such a commitment.

Does this settle the question, then? Howson is simply wrong: BCT in its

6. Though Howson and Albert consider only the deterministic case in which the relevant likelihoods are either zero or one.

most popular form does have significant inductive commitments. Modern Bayesianism is not merely an inductive framework. Yet, I will argue in the final section of this paper, there is a sense in which Howson is right. Modern Bayesianism's commitment to the likelihood lover's principle is due to its commitment to the probability coordination principle, not to its commitment to the rule of Bayesian conditionalization. The *Bayesianism* in modern Bayesianism is a framework; it is the addition of PCP to the framework that introduces the inductive commitment, and in so doing, creates an account of confirmation that can properly be called an inductive logic.

4. Probability Coordination and Induction

To get some sense of PCP's importance, ask: is Bayesian confirmation theory without PCP committed to the likelihood lover's principle? The answer is no.[7] One hypothesis may ascribe a much higher physical probability to some piece of evidence than another, but without PCP, there is nothing to stop you assigning almost any subjective conditional probabilities you like. Thus, you can set the subjective likelihood $C(e|h)$ very low for the hypothesis that assigns a high physical probability to e, and very high for the other hypothesis. Then the hypothesis that assigns the lower physical probability to the evidence will get a bigger boost from the evidence, in violation of the likelihood lover's principle.

This does not in itself show, however, that it is PCP that contains the inductive commitment to the likelihood lover's principle. It may be that the commitment is inherent in the Bayesian apparatus, but that PCP plays an indispensable role in making the commitment explicit. On this view, PCP

7. Except in the deterministic case where the hypotheses all entail either the evidence or its negation.

is like a light switch: the light does not shine unless the switch is on, but it is not the switch that powers the light.

This is, I think, the consensus view about the role that PCP plays in modern Bayesianism's commitment to the likelihood lover's principle (though few philosophers, perhaps, have any view on this matter at all). There are two reasons why it seems implausible that PCP should harbor an inductive commitment to LLP.

First, PCP has the character of a principle of direct inference, that is, a principle that tells you what to expect of the world given some statistical law. It tells you, for example, to adopt a very low subjective probability for the event of ten tosses of a coin all landing heads. But, being a principle that says something about particular events in the light of the statistical laws, it seems unlikely that it also does what an inductive commitment does, which is to say something about the statistical laws in the light of particular events. Or so you might think.

Second, and I would say more importantly, PCP's role seems to be simply one of translation. What it does is to translate the likelihoods from the language of physical probability into the language of subjective probability, that is, into the language of Bayesianism. As such, it puts the likelihoods in the right form for the application of the Bayesian apparatus, and that is all: it does not specify in what way the apparatus should be applied. That is, PCP does not tell you what to do with the likelihoods, and in particular, it does not tell you to favor hypotheses with higher likelihoods. Probability coordination is essential to modern Bayesianism, on this view, because it gives the Bayesian access to the physical likelihoods (Lewis (1980) thinks it is our *only* access); it does not, however, comment on the inductive significance of the likelihoods.

This is a very plausible line of thought, but, I have come to realize, it is entirely mistaken. Assigning a particular value to a subjective likelihood does commit you, more or less, to favoring some hypotheses over others. By directing such assignments, PCP makes inductive recommendations, in particular, recommendations in accordance with the likelihood lover's principle. (I will come back to that *more or less* shortly.)

In order to see better what kind of inductive commitments might be inherent in PCP itself, I will put Bayesian conditionalization to one side, and I will ask what kind of inductive logic you would obtain if you subscribed to PCP alone. I will call the answer the *PCP-driven logic*.

The probability coordination principle does just one thing: it dictates a value for the subjective likelihood, $C(e|h)$. But what is $C(e|h)$? It is the proportion of $C(h)$ that corresponds to $C(he)$. The rest corresponds to $C(h\neg e)$. So PCP tells you what proportion of $C(h)$ to allocate to $C(he)$, and what proportion to allocate to $C(h\neg e)$.

Now, what happens if e is observed? By anyone's lights, the part of $C(h)$ corresponding to $C(h\neg e)$ should go to zero. Thus, your subjective likelihood $C(e|h)$ for h is, in a sense, your opinion as to how much of your $C(h)$ should go and how much should stay when e is observed.

Of course, you cannot simply set $C(h\neg e)$ to zero for every h and leave it at that, or your probabilities over all the hypotheses will sum to less than one. The simplest thing to do to rectify the situation is to normalize the probabilities, that is, to multiply them all by the same factor, so that they once more add up to one. That factor will be $1/C(e)$.

The inductive procedure derived from PCP alone, then—the PCP-driven inductive logic—is as follows:

1. Assign subjective likelihoods as required by PCP.
2. Truncate: When e is observed, set $C(h\neg e)$ to zero for every h. In other words, throw away the part of the probability corresponding to $C(h\neg e)$.
3. Normalize: Multiply all probabilities by the same factor so that they again sum to 1.

The procedure is shown in figure 1. It has all the elements of an inductive logic, and, of course, it implements the likelihood lover's principle.

Now observe that the PCP-driven logic yields exactly the same changes in probability, given the observation of e, as modern Bayesianism. It seems that, in deriving the PCP-driven inductive logic, I have unwittingly adopted some Bayesian principles. More precisely, I propose, steps (2) and (3) are just a description of the operation of the Bayesian conditionalization rule. My derivation of the "PCP-driven logic" is nothing new, then; it is only modern Bayesianism derived in an unfamiliar way, taking PCP, rather than the mathematics of subjective probability, as the starting point, and so making PCP, as it should be, the focal point of, not an addendum to, the argument.

The point of the exercise is to pinpoint the source of modern Bayesianism's commitment to the likelihood lover's principle. Modern Bayesianism is, I repeat, equivalent to the PCP-driven logic: its use of PCP is equivalent to step (1) of the logic, while Bayesian conditionalization is equivalent to steps (2) and (3). Whether modern Bayesianism's commitment to the likelihood lover's principle comes from PCP or from the conditionalization rule depends, then, on whether the commitment to the likelihood lover's principle is made in step (1) or in steps (2) and (3).

Clearly, the great part of the inductive commitment to physical likelihoods is in step (1), because here it is decided in advance which hypotheses will benefit and which will suffer from the observation of any particular piece of evidence. Step (2) merely enforces the decision; step (3) preserves the outcome of the decision—the relative standing of the different hypotheses after step (2)—while restoring mathematical order to the apparatus. I conclude that PCP is the source of modern Bayesianism's commitment to the likelihood lover's principle. In addition, I conjecture that any reasonable system of inductive logic that incorporates PCP will *thereby* take on board a commitment to the likelihood lover's principle.

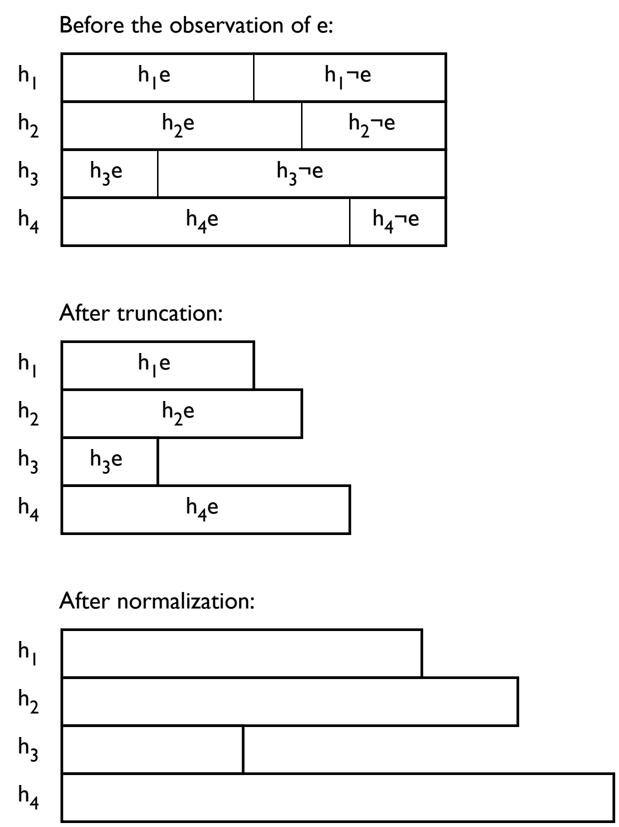

Figure 1: The PCP-driven logic in action

Observe that steps (2) and (3) on their own look very much like an inductive framework. The framework is transformed into a logic by supplying a method for apportioning $C(h)$ between $C(he)$ and $C(h\neg e)$. Modern Bayesianism commits itself to such a method, in the form of PCP, and is for this reason a true inductive logic. Bayesianism without PCP, however, appears to be a framework with almost no inductive commitment. (But only *almost* no commitment, for two reasons. First, even without PCP, Bayesianism favors hypotheses that entail the evidence over those that do not. Second, as Earman (1992, chap. 9), Kelly (1996), and others have noted, simply to adopt the apparatus of subjective probability seems to constitute a kind of bet that the world will not turn out to be a certain way.)

5. Conclusion

Bayesian confirmation theory without PCP is little more than an inductive framework. But modern Bayesianism adds PCP to the framework. This principle contains a real inductive commitment: it implements the likelihood lover's principle. If you want to know whether modern Bayesianism succeeds in justifying a certain sort of inductive behavior, then, you must ask not, as almost everyone concerned with this question until now has,

Is Bayes' rule justified?

but, with Strevens (1999),

Is the probability coordination rule justified?

References

Albert, M.: 2001, 'Bayesian Learning and Expectations Formation: Anything Goes'. In: D. Corfield and J. Williamson (eds.): *Foundations of Bayesianism*. Dordrecht: Kluwer.

Carnap, R.: 1950, *Logical Foundations of Probability*. Chicago: Chicago University Press.

Dawid, A. P.: 1982, 'The Well-Calibrated Bayesian'. *Journal of the American Statistical Association* **77**, 604–613.

Earman, J.: 1992, *Bayes or Bust?* Cambridge, MA: MIT Press.

Howson, C.: 2001, *Hume's Problem: Induction and the Justification of Belief*. Oxford: Oxford University Press.

Howson, C. and P. Urbach: 1993, *Scientific Reasoning: The Bayesian Approach*. Chicago: Open Court, 2 edition.

Jaynes, E. T.: 1983, *Papers on Probability, Statistics, and Statistical Physics*. Dordrecht: D. Reidel.

Kelly, K. T.: 1996, *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.

Keynes, J. M.: 1921, *A Treatise on Probability*. London: Macmillan.

Levi, I.: 1980, *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.

Lewis, D.: 1980, 'A Subjectivist's Guide to Objective Chance'. In: R. C. Jeffrey (ed.): *Studies in Inductive Logic and Probability*, Vol. 2. Berkeley, CA: University of California Press.

Lewis, D.: 1994, 'Humean Supervenience Debugged'. *Mind* **103**, 473–490.

Miller, D.: 1966, 'A Paradox of Information'. *British Journal for the Philosophy of Science* **17**, 59–61.

Strevens, M.: 1995, 'A Closer Look at the 'New' Principle'. *British Journal for the Philosophy of Science* **46**, 545–561.

Strevens, M.: 1999, 'Objective Probabilities as a Guide to the World'. *Philosophical Studies* **95**, 243–275.