

No Model, No Inference: A Bayesian Primer on the Grue Problem¹



Elliott Sober

1. Introduction

Stripped to its bare quintessence, the grue problem reduces to two issues. The first concerns the relationship between *generalizations* and their *instances*. If I've observed a number of emeralds before now and have found each to be green, what epistemic difference is there between the following two generalizations?

(ALLGREEN) All emeralds are green.
(ALLGRUE) All emeralds are grue.²

The second problem concerns the relationship between *predictions* and their *precedents*. If I've observed a number of emeralds before now and have found each to be green, what epistemic difference is there between the following two predictions?

(NEXTGREEN) The next emerald I examine will be green.
(NEXTGRUE) The next emerald I examine will be grue.

The phrase "epistemic difference" is vague and requires clarification. However it is understood, it is meant to rule out some obvious differences between the two hypotheses. The GREEN hypotheses are stated in a familiar vocabulary, whereas the GRUE hypotheses make use of a made-up word. Unless one is prepared to regard this difference between the hypotheses as *epistemically* relevant—as somehow relevant to what we ought to believe—this obvious difference makes no difference as far as the grue problem is concerned.³ Beyond this, it falls to any solution of the problem to make the concept of epistemic difference precise.

I have tried to state a minimum formulation of the grue problem because the *problem* is too often conflated with some of the ideas that figured in

Goodman's proposed *solution*. Goodman held that ALLGREEN is confirmed by its instances, whereas ALLGRUE is not. This goes beyond my minimal formulation of the problem, in that instance confirmation is just one concept of epistemic relevance. We should not assume at the outset that this is where the relevant difference between the GREEN and GRUE hypotheses is to be found. Goodman also asserted that the two problems I've sketched bear an important connection to each other; he maintained that ALLGREEN is confirmed by its instances only if NEXTGREEN is too. Again, the grue problem should not be burdened with this assumption; rather, this is a thesis that needs to be argued for explicitly. And finally, Goodman claimed that these evidential issues bear an important connection to the problem of distinguishing lawlike from accidental generalizations. Goodman believed that lawlike generalizations (like ALLGREEN) are confirmed by their instances, whereas accidental generalizations (like ALLGRUE) are not.⁴ Having separated these Goodmanian theses from the problem they are meant to address, I will argue that each is incorrect and can be seen to be so by giving the grue problem a probabilistic representation.

During the time that Goodman wrote about the grue problem, philosophers often focused on the problem of *qualitative confirmation*. Just as Goodman had his emeralds, Hempel (1965) had his ravens. Hempel's problem was not to measure *the degree of confirmation* that observing red shoes provides for the generalization "all ravens are black." Rather, he wanted to say whether such observations provide any confirmation at all. Now, quite apart from the merits of what Goodman and Hempel said about emeralds and ravens, it seems undeniable that a theory of qualitative confirmation should be embeddable in a theory of *quantitative* confirmation. We need the notion of *degree of confirmation*, not just the dichotomous division of *confirmed* versus *not confirmed*. The scientific community employs probability concepts in its understanding of this quantitative notion. So a natural place to begin discussion of the grue problem is with the concept of probability.

2. Bayesianism—The Basics

Those who agree that probability is a useful concept for explicating epistemic concepts nonetheless disagree about how it should be deployed. Bayesianism constitutes an influential school in this debate, but it is not the only game in town.⁵ I'll mention later on a standard criticism of Bayesianism. But for now, Bayesianism is a perfectly sensible place to begin. For many inference problems, the Bayesian solution is not at all controversial, in that it happens to coincide with the verdicts, if not the exact

reasoning, of other approaches. In addition, the broad lessons I will draw from my Bayesian analysis are *robust*; they are very much in harmony with the conclusions that would be generated by non-Bayesian positions.

Bayesians think that observations confer probabilities on hypotheses and that the mathematical idea of probability is the right way to measure the epistemic property of plausibility. Bayes's theorem describes how the *posterior probability* of hypothesis H —the probability it has in the light of the observation O —is a function of three other quantities:

$$P(H/O) = P(O/H)P(H)/P(O).$$

$P(O/H)$ is called the likelihood of H . It describes the probability that the hypothesis H confers on the observations. Don't confuse $P(O/H)$ with $P(H/O)$; the likelihood of H and the probability of H are different. $P(H)$ is termed the *prior* probability of H , meaning the probability the hypothesis has *before* the observation O is made.

The grue problem, whether it is understood in terms of generalizations or in terms of predictions, involves *comparing* two hypotheses in the light of a body of observations. By bringing together two applications of Bayes's theorem, we can derive a comparative principle of the kind required:

$$(S) \quad P(H_1/O) > P(H_2/O) \text{ if and only if } P(O/H_1)P(H_1) > P(O/H_2)P(H_2).$$

If H_1 is to have a higher posterior probability than H_2 , this must be because H_1 has the higher likelihood or the higher prior probability (or both). So if we are to find a difference between the GREEN and the GRUE hypotheses in this Bayesian format, we know that there are exactly two places to look.

Principle (S) is a *synchronic* principle. It does not describe how much *change* the observation O engenders in the probabilities of the hypotheses; it simply describes what it takes for the one to have a higher value than the other, after the observations are obtained. However, the *diachronic* issue is also worth considering. If $P(H/O)$ represents the plausibility that H has *after* O is found to be true, and $P(H)$ is the plausibility that H possesses *before* that event, then it is natural to describe the *change* in plausibility that H experiences across this change as follows:

$$\begin{aligned} O \text{ confirms } H & \text{ if and only if } P(H/O) > P(H). \\ O \text{ disconfirms } H & \text{ if and only if } P(H/O) < P(H). \end{aligned}$$

Notice that O can confirm H even though $P(H/O)$ is quite low; merely let $P(H) = 0.000001$ and $P(H/O) = 0.01$. And O can disconfirm H even

though H remains quite probable in the light of O ; merely let $P(H) = 0.95$ and $P(H/O) = 0.94$.

Two applications of these definitions yield the following diachronic principle:

- (D) O confirms H_1 more than O confirms H_2 if and only if $P(H_1/O) - P(H_1) > P(H_2/O) - P(H_2)$.

I take the *difference* between the posterior and prior probabilities, rather than the *ratio* between them, to represent degree of confirmation. Ellery Eells has suggested to me the following argument for this choice. Suppose that

$$P(H_2/O) = 0.9 \quad P(H_1) = 0.1$$

and that

$$P(H_2/O) = 0.001 \quad P(H_2) = 0.0001.$$

If degree of confirmation is measured by taking differences, H_1 is confirmed by O more than H_2 is. But if degree of confirmation is measured by the ratio of posterior to prior, the reverse is true. Surely a jump from 0.1 to 0.9 reflects a larger change in plausibility than a jump from 0.0001 to 0.001. In any event the main conclusions I will argue for in what follows do not turn on my choice of measure.

Although Goodman formulated his puzzle as one about *confirmation*, there is no reason to restrict our attention to the diachronic issue. We want to know if observing green emeralds raises the probabilities of the GREEN hypotheses more than such observations raise the probabilities of the GRUE hypotheses. But we also would like to know whether the GREEN hypotheses are more probable, in the light of these observations, than the GRUE hypotheses are. So there are four cases to consider—the generalization and the prediction problems each need to be considered both diachronically and synchronically.

Even though the ideal of minimalism has guided me in my description of the grue problem, I have to admit that I introduced a substantive choice in my formulation. This involves the way I have described the evidence. I imagine that our evidence was obtained by examining some emeralds and finding out what color they exhibit. This search strategy differs from that of sampling the universe at random and noting whether the things we come up with are emeralds and what their colors are. With fairly modest assumptions, it turns out that sampling the emeralds is a better strategy than sampling the world at random (if the goal is to test either of the two generalizations described above).⁶ I assume in what follows that this is how our evidence was obtained. But once again, many of the broad lessons I will draw from my analysis do not turn on the details of this assumption.

3. Generalization—The Diachronic Question

Let's begin with a simple point about Bayesianism. If H deductively entails O and if O was not certain to be true before the observation was made, then O confirms H . This is true because Bayes's theorem can be rewritten as follows:

$$P(H/O)/P(H) = P(O/H)/P(O).$$

O confirms H precisely when the left-hand side is greater than 1, which must be so, if $P(O/H) = 1$ and $P(O) < 1$.

What does this mean about the grue problem? Here we must tread carefully. Let us begin with the formulation that focuses on generalizations. Notice that ALLGREEN and ALLGRUE both deductively imply that the emeralds examined before now are green.⁷ If it was not a certainty beforehand that those emeralds should have turned out to be green, we must conclude that both hypotheses are confirmed by the past observations. Where the proposition E says that the sampled past emeralds are green, the relevant facts are

$$P(\text{ALLGREEN}/E)/P(\text{ALLGREEN}) = P(E/\text{ALLGREEN})/P(E) > 1$$

and

$$P(\text{ALLGRUE}/E)/P(\text{ALLGRUE}) = P(E/\text{ALLGRUE})/P(E) > 1.$$

Both generalizations are confirmed by their instances if this simply means that each has its probability increased by the past observations.

But now let us consider the degree of confirmation that each generalization experiences. I begin by noting that

$$P(\text{ALLGREEN}/E) - P(\text{ALLGREEN}) > P(\text{ALLGRUE}/E) - P(\text{ALLGRUE})$$

if and only if

$$P(E/\text{ALLGREEN})P(\text{ALLGREEN})/P(E) - P(\text{ALLGREEN}) > P(E/\text{ALLGRUE})P(\text{ALLGRUE})/P(E) - P(\text{ALLGRUE}).$$

This latter inequality is true precisely when

$$P(\text{ALLGREEN})(1/P(E) - 1) > P(\text{ALLGRUE})(1/P(E) - 1),$$

which simplifies to $P(\text{ALLGREEN}) > P(\text{ALLGRUE})$, if $P(E) < 1$. In other words, if we assume that the observations were not certain beforehand and that each generalization implies the observations, then ALLGREEN is confirmed more than ALLGRUE if and only if the former has the higher prior probability.

What would make it plausible to assign ALLGREEN the higher prior probability? Rather than addressing that question here, I postpone it until the next section.

4. Generalization—The Synchronic Question

As noted before, if we are sampling from the population of emeralds, then the generalizations ALLGREEN and ALLGRUE each deductively entail that the items sampled before now were green. Since the two hypotheses both have likelihoods of unity, principle (S) entails that a difference in posterior probability must be due entirely to a difference in prior:

If $P(O/H_1) = P(O/H_2)$, then $P(H_1/O) > P(H_2/O)$ if and only if $P(H_1) > P(H_2)$.

So ALLGREEN is more probable than ALLGRUE, given the evidence E , precisely when the former hypothesis has the higher prior.

What could justify the belief that ALLGREEN has a higher prior probability than ALLGRUE? It is at this point in the story that I must confess my anti-Bayesian sympathies. If prior probabilities are to be objective, I do not see how they can be assigned a priori. And if they are merely subjective—simply indicating the degree of belief of some agent—then I don't see that they have any epistemic relevance to this problem. One does not show that ALLGREEN is more plausible than ALLGRUE simply by giving voice to the autobiographical remark that one finds the former more plausible a priori than the latter.

Furthermore, I do not see how these hypotheses can have (objective) probabilities unless they describe possible outcomes of a chance process (Edwards 1972). I do not know what it means to say that Newton's law of gravitation or Darwin's theory of evolution has an objective probability. They were not made true by God's reaching into an urn that contained slips of paper on which candidate laws were inscribed. So if we can't specify a chance process that produces a coloration pattern for emeralds, I don't know what it means to assign ALLGREEN and ALLGRUE objective prior probabilities.

Having said that I find both objective and subjective Bayesianism

unattractive as general doctrines, I do not think we must concede that the problems at hand are insoluble. If we find ALLGREEN more plausible than ALLGRUE, this is because we hold various substantive, if hard to articulate, theories about the world. Perhaps we expect emeralds to be alike in color because we think that they are alike in physical structure, and we believe that color supervenes on physical structure. Of course, these convictions involve assumptions about the future. But there is no escaping such commitments; based solely on our experience of the past, the generalizations cannot be shown to differ in their probabilities.

5. Prediction—The Diachronic Question

Let us now shift from the issue of generalization to the issue of prediction. NEXTGREEN does *not* deductively imply that the emeralds examined before now have been green, and neither does NEXTGRUE. So the arguments that solve the problems associated with the generalizations do not apply to the problems about prediction. Is there some other argument that forces the same conclusion? Or are generalizations and predictions not as tightly coupled epistemologically as Goodman and many others have thought?

To answer this question, let us leave the strange and wonderful world of grue behind for a moment, and consider a rather more mundane inductive problem. Imagine an urn that is filled with a thousand balls by drawing from a source whose composition is known. Suppose the source contains 50 percent red balls and 50 percent green balls. By random sampling from the source, the urn is filled with a thousand balls. The problem is to sample from the urn (with replacement) and to draw two inferences based on the sample obtained. The first inference is to be a generalization concerning the composition of the whole urn. The second is to be a prediction concerning the color of balls that will be sampled in the future.

Since we know that the urn was composed by draws from the source, we can assign prior probabilities to each of the possible compositions, from 1000 red and 0 green to 0 red and 1000 green. When we sample from the urn, we can use Bayes's theorem to compute the posterior probability of the various hypotheses about the urn's composition. Suppose I take 250 draws from the urn and find that each ball I sample is green. These observations make the hypothesis that all the balls in the urn are green more probable than it was initially. Just as was true for ALLGREEN and ALLGRUE, the generalization is confirmed when a prediction deduced from the generalization comes true, provided that the prediction was not certain beforehand.

But now let us consider the problem of predicting what the next ball will

be like, given information about the character of the sample. In this case, the probability that the next ball will be green is 0.5, regardless of what the previously sampled balls were like. Observing 250 green balls does *not* raise the probability that the next ball drawn from the urn will be green, even though the observations do raise the probability that all the balls are green. How can we make sense of the fact that generalization and prediction part ways in this example?

We can grasp the general point by considering the simple case in which we make just two draws (with replacement) from the urn. There now are four equiprobable sequences of green (*G*) and red (*R*) draws—*GG*, *GR*, *RG*, and *RR*. If the first ball sampled is green, then the probability of *GG* increases from 0.25 to 0.5. However, after this first observation, the probability that the second draw will be green is still 0.5, just as it was before any ball was drawn. I conclude that confirming a generalization and confirming a claim about the next instance are not always as intimately connected as Goodman suggests.

This example illustrates another defect in what Goodman said about the grue problem. He claimed that a generalization is confirmable by its instances only if it is lawlike. This latter concept, whatever else it might mean, entails that the generalization “supports” a counterfactual. If “All *Xs* are *Ys*” is lawlike, then if it is true, so is the statement “If *a* were an *X*, then *a* would be a *Y*.” In the example just given, the generalization “All balls in the urn are green” is confirmed by its instances. Yet knowledge of the process whereby the contents of the urn were assembled guarantees that this generalization, if true, is only accidentally so. It is a mere fluke if all the balls in the urn happen to be green. There is nothing about being a ball in this urn that makes something green, nor is it true that a ball would not have been put into the urn unless it were green. “If this tennis ball were in the urn, then it would be green” is as false as any counterfactual can get.⁸

When questions of lawlikeness are considered, it makes all the difference in the world whether the mechanism whereby the population is assembled is known in advance or is inferred in the process of sampling. If I sample balls from the urn and find that all are green, and I have no idea how the urn was formed, the suspicion naturally arises that the homogeneous character of my sample is not an accident. However, the fact that nomological connections are reasonably suspected in such cases does not show that lawlikeness is a presupposition of instance confirmation. Simply replace prior ignorance of process with a substantive process assumption (of the kind just sketched for the urn problem), and the composition of a population known to be fortuitously assembled can be confirmed by random sampling.⁹

The urn example establishes that confirming a generalization does not

require that one confirm a prediction about the next instance. Sampling from the population of emeralds, it is inevitable (given the modest assumption that it was not a certainty that the sampled emeralds would turn out to be green) that observing green emeralds before now should raise the probability that all emeralds are green (ditto for grue). But no such inevitability attaches to the prediction that the next emerald examined will be green (or grue). Having separated these two problems, let us now explore the prediction problem on its own.

We have observed emeralds before now and found each to be green. What does it take for that body of observation (*E*) to raise the probability that the next emerald I examine will be green? A useful representation of when this is true is provided by the following:

$$P(\text{NEXTGREEN}/E) > P(\text{NEXTGREEN}) \text{ if and only if } \\ P(\text{NEXTGREEN} \ \& \ E) > P(\text{NEXTGREEN})P(E).$$

The right-hand side of this biconditional says that the *covariance* of *NEXTGREEN* and *E* is positive. The covariance of *A* and *B* is defined as $\text{Cov}(A, B) = P(A \ \& \ B) - P(A)P(B)$. When *A* and *B* are independent of each other, $\text{Cov}(A, B) = 0$. Positive covariance means positive association.

So the GREEN prediction is confirmed by the evidence only if the prediction and the evidence exhibit positive covariance. The 2 × 2 table below represents this constraint on the confirmation of *NEXTGREEN*. In Table 1, $P(E) = p$ and $P(\text{NEXTGREEN}) = q$. A positive covariance means that $c > 0$.

As noted earlier in connection with the urn example, it is not inevitable that past precedents should confirm a prediction. The information presented in that example concerning how the urn was filled made all the difference. I suggest that a similar answer be given in connection with this problem about *NEXTGREEN*. If *NEXTGREEN* is confirmed by *E*, this is because some sort of process induced a correlation (a positive covariance)

		(<i>q</i>) NEXTGREEN	(1 - <i>q</i>) NOT-NEXTGREEN
(<i>p</i>)	<i>E</i>	$pq + c$	$p(1 - q) - c$
(1 - <i>p</i>)	NOT- <i>E</i>	$(1 - p)q - c$	$(1 - p)(1 - q) + c$

Table 1

between past and future emeralds with respect to their color.

What sort of process might this be? In this respect Goodman's example is a bit unfortunate, since emeralds, I gather, are standardly said to be green *by definition*. But ignoring this wrinkle in Goodman's example, a natural suggestion is that the relevant process assumption is that emeralds had their color determined *as a group*. If emeralds share the same microstructure, and if microstructure determines color, then one has the basis for expecting, before even one emerald is examined, that emeralds will be alike in color.

But suppose that no such process assumption is available. If we know nothing about the process by which emeralds receive their colors (or grulers), how are we to decide whether past precedents confirm a prediction? The answer, I think, is that *we cannot*. As Hume argued, a description of the past, in and of itself, offers no guidance whatever as to what the future will be like. Here I don't mean just that we can't *deduce* what the future will be like from a description that is solely about the past. The Humean point is more profound: we can't even infer what the future will *probably* be like, based solely on a description of the past. Nor can we say whether past precedents confirm a prediction about the future unless we are willing to make assumptions concerning how past and future are related (Sober 1988b).

I so far have explored what it takes for the prediction NEXTGREEN to be confirmed by the observation *E*. What can be said of the relationship of NEXTGRUE to the same observation? If "grue" just meant *green before or not green after*, then we could conclude that NEXTGREEN is confirmed just in case NEXTGRUE is disconfirmed. But the usual definitions of "grue" do not permit this simple conclusion to be drawn. It is possible for *E* to confirm both NEXTGREEN and NEXTGRUE. No first principle rules out the assumption that the past observations raise the probability of both predictions (even though they are incompatible with each other).

Nonetheless, we have obtained a *necessary* (but not a sufficient) condition for NEXTGREEN to be confirmed more than NEXTGRUE; for this to be so, NEXTGREEN must be confirmed, which requires that $c > 0$. And if we assume that the next emerald will be either green or blue, then $c > 0$ is both necessary and sufficient for NEXTGREEN to be confirmed and NEXTGRUE to be disconfirmed.

6. Prediction—The Synchronic Question

Given that the emeralds observed before now have all been green, is it more probable that the next emerald will be green or that it will be grue? If we reformulate this question a little, we can use the 2×2 table described before to identify an assumption on which this difference in probabilities depends:

$$P(\text{NEXTGREEN}/E) > P(\text{NOT-NEXTGREEN}/E)$$

if and only if $[pq + c]/p > [p(1 - q) - c]/p$
 if and only if $c > p(1 - 2q)/2$.

Here we have shifted the problem to asking whether past observations make it more probable that the next emerald will be green or that it will be *not* green. Since the probability of NOT-NEXTGREEN cannot be less than the probability of NEXTGRUE, the biconditional describes a necessary condition for $P(\text{NEXTGREEN}/E) > P(\text{NEXTGRUE}/E)$. And as before, if we assume that the next emerald will either be green or blue, then the condition cited is both necessary and sufficient.

I noted in connection with the diachronic prediction problem that we can't assume *a priori* that $c > 0$. The point of interest here is that this assumption does not *suffice* for the GREEN prediction to be more probable than the GRUE prediction. Here we see a difference between the synchronic and the diachronic versions of the prediction problem.

7. Summary

Let's take stock. Inferring the color of the next emerald and inferring what all emeralds are like are different inference problems. Prediction and generalization present different issues. Likewise how much an observation boosts the probability of a hypothesis is a different question from how high that probability actually becomes. The diachronic and synchronic issues need to be separated. The results obtained from applying this pair of dichotomies to the grue problem are summarized in Table 2. In each case, we can assert that an epistemic asymmetry obtains between the GREEN and the GRUE hypothesis only if we are prepared to make a substantive assumption about the way the world is. What is more, the assumptions change as we shift from problem to problem.

8. Concluding Comments

With these results in hand, it is worth stepping back for a moment to reflect on the nature of the grue problem and on what kind of solution we can hope to attain.

Once again, a comparison with Hempel's raven problem is instructive. In a paper of breathtaking brevity, Good (1967) showed that observing a black raven (sampled at random from the world at large) can actually *disconfirm* the generalization that all ravens are black, provided one adopts a few simple (if implausible) empirical assumptions about the inference problem. Hempel (1968) replied that Good's argument misconstrued

	Diachronic	Synchronic
Generalization	$P(\text{ALLGREEN}/E) - P(\text{ALLGREEN}) > P(\text{ALLGRUE}/E) - P(\text{ALLGRUE})$ if and only if $P(\text{ALLGREEN}) > P(\text{ALLGRUE}) \text{ and } P(E) < 1.$	$P(\text{ALLGREEN}/E) > P(\text{ALLGRUE}/E)$ if and only if $P(\text{ALLGREEN}) > P(\text{ALLGRUE}).$
Prediction	If the next emerald will be either green or blue, then $P(\text{NEXTGREEN}/E) - P(\text{NEXTGREEN}) > P(\text{NEXTGRUE}/E) - P(\text{NEXTGRUE})$ if and only if $c > 0$.	If the next emerald will be either green or blue, then $P(\text{NEXTGREEN}/E) > P(\text{NEXTGRUE}/E)$ if and only if $c > p(1 - 2q)/2.$

Table 2

the problem that he, Hempel, had wanted to pose. Hempel (1965) had indicated that he was interested in exploring the relationship of observation to hypothesis within a “theoretically barren” background context. Assuming nothing at all about the world, the question is whether black ravens and red shoes both confirm the generalization that all ravens are black.

Good took the view, and so do I, that almost nothing can be said about confirmation in a background context of this sort. In all four of the problems surveyed above, an epistemic asymmetry between the GREEN hypothesis and the GRUE hypothesis is possible. But notice what the asymmetries in these cases depend upon: for GREEN to be more probable than GRUE, or for GREEN to receive a greater boost in probability than GRUE does, *empirical assumptions must be made that go beyond the testimony of past observation*. It isn’t reason alone (or “the scientific method”) that induces an asymmetry here, but substantive assumptions about the way the world is.

I therefore think it is misleading, at best, to claim that the GREEN hypothesis is preferable to the GRUE hypothesis on the ground that the former is “simpler.” This appeal to simplicity gives the impression that the simplicity of a hypothesis is a reason to think that it is true. Although many

philosophers believe that appeal to such “extra-empirical” virtues is part of what it means to do science, I do not. The austere framework of Bayesianism, and of other probabilistic epistemologies, accords no irreducible role to simplicity, unification, non-*ad hocness*, and so on.¹⁰ What this means is that simplicity never provides an irreducible justification of any hypothesis. To be sure, it is sometimes true that the simpler theory is more probable, or more likely, or more strongly confirmed by a body of data; however, the simpler theory never has these properties *because* it is simpler (Sober 1988b, 1990). In the present context, it does no harm to admit that the GREEN hypothesis is “simpler” than the GRUE hypothesis. But if this difference is to count as epistemically relevant, it will be necessary to appeal to empirical matters of fact of the sort described in the above table. Once these empirical assumptions are made explicit, any further mention of simplicity will be quite unnecessary.

For me, the fundamental lesson of the grue problem is that empirical assumptions that go beyond the content of past observations are needed to establish an epistemic asymmetry between GREEN and GRUE. Whereas philosophers often formulate this point by appealing to the need for “auxiliary assumptions,” scientists of a statistical bent often stress the importance of specifying a “model” of the relation of data to the various hypotheses under test. Without assumptions of this sort, the data cannot be interpreted. The slogan for scientists is: *NO MODEL, NO INFERENCE*. This entirely familiar point from the practice of science should not be forgotten when we investigate the theory of that practice.

If this is the right lesson to draw from the grue problem, we can reach an assessment of the solution to the problem that Goodman proposed—his theory of entrenchment. A predicate becomes entrenched when people use it to formulate predictions and generalizations. It has always been a mystery to me why the fact that people use a predicate should have any epistemic relevance. Why should our use of a predicate be evidence that this or that hypothesis is true? This naive question is sometimes answered with the response that the “new” riddle of induction involves describing our inductive practices, not trying to justify them. I have my doubts about this descriptive claim as well. Is it really so obvious that human inference makes think a hypothesis with unentrenched predicates is less plausible than a hypothesis with entrenched predicates, all else being equal? But this reply to one side, I hope it is clear why Goodman’s theory is the wrong *kind* of theory, at least if one is interested in normative questions of evidence and confirmation. To describe how well entrenched a predicate *now* is involves describing *past* events only. No such description can suffice to establish an epistemic asymmetry between a GREEN hypothesis and a GRUE hypothesis.

I suspect that many philosophers who may have been skeptical of the

details of Goodman's entrenchment theory nonetheless thought that Goodman was at least looking for the right *kind* of asymmetry; they implicitly assumed that a solution to the problem could be found in some fact established solely by past experience and a priori reasoning, without recourse to induction itself. To rest an asymmetry between GREEN and GRUE on an assumption about the future was standardly said to "beg the question." But what question are we thereby begging? It is the following ill-formed question: *Which beliefs based just on past experience and on a priori reasoning suffice to show that the GREEN hypothesis is more reasonable than the GRUE hypothesis?* To me, this question resembles another: *Which types of butter are capable of cutting a diamond?*

University of Wisconsin, Madison

NOTES

1. I thank Ellery Eells, Malcolm Forster, and Douglas Stalker for helpful comments on an earlier draft of this paper. I also am grateful to the editors of the *Philosophical Review* for granting me permission to reprint here a few paragraphs from my article (Sober 1988a).

2. Since the year 2000 is fast approaching, I will define "grue" so that the grue problem retains its timeless immediacy: An object is said to be grue at a given time precisely when it is either green and the time is before now, or it is blue and the time is not before now. As has become somewhat customary, I delete the concept of "being examined" from Goodman's original definition.

3. Vocabulary differences between the two hypotheses cannot make any epistemic difference, if a *principle of logical equivalence* is correct. As Goodman pointed out in his exchange with Carnap, once "bleen" is defined in tandem with grue (bleen = blue before or green after), the GRUE hypotheses can be reformulated in familiar vocabulary and the GREEN hypotheses can be expressed by using the made-up words.

4. Here is a passage from Goodman (1965, 73) in which these three theses are asserted:

That a given piece of copper conducts electricity increases the credibility of statements asserting that other pieces of copper conduct electricity, and thus confirms the hypothesis that all copper conducts electricity. But the fact that a given man now in this room is a third son does not increase the credibility of statements asserting that other men now in this room are third sons, and so does not confirm the hypothesis that all men now in this room are third sons . . . The difference is that in the former case the hypothesis is a *lawlike* statement; while in the latter case, the hypothesis is a merely contingent or accidental generality.

Only a statement that is *lawlike*—regardless of its truth or falsity or its scientific importance—is capable of receiving confirmation from an instance of it; accidental statements are not.

5. Two alternatives should be mentioned here. First, there is likelihoodism, according to which hypotheses are evaluated solely in terms of the probabilities they confer on observations (Edwards 1972). Second, there is the approach of Akaike (1973) and his school; see Forster and Sober (forthcoming) for discussion.

6. Discussion of the raven problem from a Bayesian point of view has made it abundantly clear why looking at ravens and seeing whether they are black is a better strategy for testing "All ravens are black" than the strategy of looking at nonblack things and seeing if they are nonravens (or, for that matter, the strategy of looking at objects drawn from the whole universe and seeing if they are consistent with the hypothesis). See, for example, Chihara (1981), Horwich (1982), Eells (1982), Howson and Urbach (1989), and Earman (1992).

7. Of course, ALLGREEN does not deductively imply that an object sampled at random from the whole universe will be both an emerald and green. Recall that I am assuming that the sampling takes place within a restricted universe; it is part of the sampling problem that the population of objects from which the sample is drawn is composed entirely of emeralds. The relevant fact here is that for each object a sampled before now from this population, $P(a \text{ is green} / \text{ALLGREEN} \ \& \ a \text{ is an emerald}) = P(a \text{ is green} / \text{ALLGRUE} \ \& \ a \text{ is an emerald}) = 1$.

8. I discuss what other philosophers have said about Goodman's proposed connection between confirmation and lawlikeness in Sober (1988a).

9. Many sciences (e.g., population biology) test generalizations by drawing samples from populations that are known to be fortuitously assembled. This undertaking would be impossible if only lawlike statements could be confirmed by their instances. It is worth remembering that Goodman's proposal was advanced during a period in which philosophers often equated science with physics and physics with the search for physical laws.

10. Of course, one might take this fact to be a *reductio* of such probabilistic epistemologies; one might argue that they are mistaken precisely because they fail to accord an irreducible importance to these "epistemic virtues." An assessment of this suggestion can be developed only by attending closely to the dialectics of theory choice in a variety of scientific controversies.

REFERENCES

- Akaike, H. (1973). "Information theory and an extension of the maximum likelihood principle." In B. Petrov and F. Csaki, eds., *Second International Symposium on Information Theory*, Akademiai Kiado, 267–81.
- Chihara, C. (1981). "Quine and the confirmational paradoxes." In P. French, H.

- Wettstein, and T. Uehling, eds., *Midwest Studies in Philosophy*, vol. 6, University of Minnesota Press, 425–452.
- Earman, J. (1992). *Bayes or Bust*. MIT Press.
- Edwards, A. (1972). *Likelihood*. Cambridge University Press.
- Eells, E. (1982). *Rational Decision and Causality*. Cambridge University Press.
- Forster, M. and Sober, E. (forthcoming). "How to tell when simpler, more unified, or less *ad hoc* theories will provide more accurate predictions."
- Good, I. (1967). "The white shoe is a red herring." *British Journal for the Philosophy of Science* 17, 322.
- Goodman, N. (1965). *Fact, Fiction and Forecast*. Bobbs-Merrill.
- Hempel, C. (1965). "Studies in the logic of confirmation." In *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. Free Press.
- Hempel, C. (1968). "The white shoe—no red herring." *British Journal for the Philosophy of Science* 18, 239–40.
- Horwich, P. (1982). *Probability and Evidence*. Cambridge University Press.
- Howson, C. and Urbach, P. (1989). *Scientific Reasoning: The Bayesian Approach*. Open Court.
- Sober, E. (1988a). "Confirmation and law-likeness." *Philosophical Review* 97, 93–98.
- Sober, E. (1988b). *Reconstructing the Past: Parsimony, Evolution, and Inference*. MIT Press.
- Sober, E. (1990). "Let's razor Ockham's razor." In D. Knowles, ed., *Explanation and Its Limits*. Cambridge University Press. 73–94.

Bayesian Projectibility

Brian Skyrms



Undoubtedly we do make predictions by projecting the patterns of the past into the future, but in selecting the patterns we project from among all those that the past exhibits, we use practical criteria that so far seem to have escaped discovery and formulation.

Nelson Goodman, "A Query on Confirmation" (1946)

I. Introduction

In 1946 Nelson Goodman raised the problem of the projectibility of hypotheses in a note addressing the confirmation theory of Rudolf Carnap. He later gave a sensational illustration: the hypothesis "All emeralds are grue," which came to be discussed as the Goodman Paradox. The predicate "grue" was so manifestly pathological many were led into thinking that the problem was to find some criterion which would exclude similar pathology from inductive reasoning.

Goodman sees clearly that the problem of projectibility is much more general, and that judgments of projectibility must be central to any adequate theory of confirmation. He believes that a theory of projectibility should be a pragmatic theory, and sketches the beginnings of such a theory in the last chapter of *Fact, Fiction and Forecast*.

I believe that the broad outlines of Goodman's approach to a theory of projectibility are just right. The theory will not attempt to tell one in a vacuum which predicates are projectible. Rather it will explain how present judgments of projectibility should be based on past judgments of projectibility together with the results of past projections. The circularity involved in such a theory is not to be viewed as vicious. The theory can still inform our understanding of projectibility. And a theory with this sort of "virtuous circularity" is really the best that can be expected. However, the implementation of Goodman's program for a theory of projectibility does not seem very far advanced. There is, however, a preexisting pragmatic framework which offers precise tools for addressing the question: the theory of personal probability. How is projectibility represented within this framework? This question leads straight to central concepts of Bayesian statistics.