
Toward a Rational Reconstruction of Design Inferences

TIMOTHY MCGREW
*Department of Philosophy
Western Michigan University
Kalamazoo, Michigan*

*From long habit the train of thoughts ran so swiftly through my mind
that I arrived at the conclusion without being conscious of intermedi-
ate steps. There were such steps, however.*

Sir Arthur Conan Doyle, *A Study in Scarlet*

Detecting the deliberate operation of intelligent agents in the environment is an activity so natural and pervasive that we may not be inclined, pre-reflectively, to classify it as an inference at all. The analogy with visual perception is initially very attractive. Just as we spontaneously discriminate among physical objects without any overt sense of deliberation, so we find ourselves sorting out signals from noise, extracting artifacts from archaeological digs, and mentally separating mechanical contrivances from their environmental contexts without weighty deliberation. It is a poor handyman who needs the help of philosophy to distinguish the lawnmower from the lawn.

But the detection of design also arises in contexts where, for a variety of reasons, the clarity of our vision shades off toward opacity. No doubt, Spenser wrote the *Amoretti*, but did he deliberately space some of the sonnets so as to reflect the liturgical year, or is their spacing a coincidence? Does the feathered tail really go with the dinosaur body in this fossil, or are we being subjected to a clever hoax? Shall we classify the untimely demise of Birdy Edwards as a homicide or as an accidental death? No reasonable person doubts that in such cases we do well to pay close attention to a wealth of details that might not be apparent to the casual observer. Manuscript comparisons, radioactive dating, the patient sifting of forensic details—the strategies we employ to disentangle design from chance are legion, and they make little sense if the detection of design is essentially a simple and unreflective matter of perception.

In the clear-cut cases, we often have little or no difficulty achieving consensus regarding design or chance: our common intuitions are robust. And even in the most complex cases where we are at a loss to explain the phenomena at hand, we are not wholly without rational guidance. Often we are able to see quite clearly what sort of evidence would, if only we could acquire it, strengthen or weaken the case for design. We have strong and stable intuitions regarding which aspects of our background knowledge are crucial and what sorts of new information would be relevant or irrelevant.

In view of the sensitivity of our judgments to available evidence, it makes sense to try to reconstruct our detection of the deliberate activities of intelligent agents as an inference—to analyze both our prereflective practices and our more painstaking investigations with a view to making explicit their underlying rational structure. But before we embark on such a rational reconstruction, we need both a clear sense of what the project is about and some criteria by which to evaluate rival proposals. In the first section of this paper, I discuss the potential tension between the rationality and the verisimilitude of our reconstructions and suggest that despite the difficulties inherent in this project the game is very much worth the candle. In the second section, I propose seven criteria and argue that they are plausible constraints on any rational reconstruction of design inferences. In the next three sections, I examine some current proposals regarding the methodology of design detection in light of these criteria, and in the final section, I draw some conclusions from this examination.

The Very Idea of Rational Reconstruction

To engage in rational reconstruction is to assume that we have something to reconstruct and some basis on which to proceed. In the case at hand, the basis comprises a large set of instances in which there are shared intuitions regarding the rationality of concluding design. The assumption behind the project is that our widespread agreement indicates the existence of an underlying rational structure, a structure perhaps glimpsed in our more careful deliberations but present in the background even when our judgments of design seem spontaneous.

One step in rational reconstruction of design inferences is a sort of inference to the best explanation. We attempt to account for the fact of our common intuitions—a fact that seems surprising given the relatively casual way in which we make many of our design judgments—inferring that there exists a rational structure linking our evidence to our judgments and that in making those judgments we are tacitly responding to that structure. Having taken this step, we are in a position to examine particular proposals regarding the nature of that underlying structure.

There are rival explanatory projects that would not count as *rational* reconstructions at all: explanations for our common intuitions in terms, for example, of social class or historical situation. The question is not whether such explanations might be rationally supported (clearly they might) but whether, if we were to accept them at face value as the true and proper account for our shared intuitions, we could still consider those intuitions themselves to be rational in anything more than a Pickwickian sense. To count as a rational reconstruction, our explanation must entail that our judgments meet certain standards of rationality.

The project is not guaranteed to be successful, either in whole or in part. We might conclude that we were wrong in some of the initial judgments, perhaps because they fail to accord with a reconstruction that is, on its own terms, wholly compelling. We might even decide that we were wrong to ascribe rationality to any of our design judgments. Of course, our mere failure to discover an adequate rational reconstruction would count only as a very weak reason to despair. But the search for a rational reconstruction could, in principle, unearth devastating skeptical arguments that would provide a powerful reason to believe that there is no way to complete the project.

A second possibility of failure arises out of the well-known fallibility of human reasoners. Psychological research has shown past reasonable doubting that we humans are liable, particularly in our artless hypothesizing, to make systematically poor judgments—sometimes spectacularly so, and sometimes in aggregate.¹ We are also limited in our cognitive resources, not only in obvious and predictable ways such as limited memory and low computational speed but also sometimes in unpredictable ways such as memory loss or inadvertence.

A rational reconstruction should not display our avoidable cognitive faults. But an excessively rigorous reconstruction may demand more than we can give, leaving us with the unhappy conclusion that although *someone* might reasonably conclude design on the basis of such evidence as is in the theory available to us, our own conclusions are rarely or never well founded. There is, then, a potential tension between the rationality of our reconstructions on the one hand and their verisimilitude—their fidelity to our actual underlying cognitive processes, even when those processes are operating at their best—on the other.

Despite this difficulty, there are reasons for optimism. In acknowledging the fallibility of our shared intuitions, we have already admitted that we may have to withdraw some of our initial judgments. Provided that we can

¹ The classic collection of research papers on this topic is Kahneman, Slovic, and Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge: Cambridge University Press, 1982). Robyn Dawes provides a compact survey of some of the salient systematic biases in human cognition from a Bayesian standpoint in *Rational Choice in an Uncertain World* (New York: Harcourt Brace, 1988).

do so in a principled way, this is no bar to the project: we can accommodate a certain amount of human error by weeding out the poor judgments. And many of the sources of individual fallibility are swamped out when we work with a robust set of shared judgments.

There are also reasons for perseverance. A good reconstruction may enable us to give a better, more precise case for judgments of design because it draws our attention to factors we are systematically prone to overlook. We start, naturally enough, with our base of common intuitions; and our ability to give a plausible account for a large proportion of those is the touchstone of success for the verisimilitude of the reconstruction. But the rationality of the reconstruction is an intrinsic feature. It is not dependent on the blind faith that we must at least fairly often arrive at our beliefs as we ought; rather, in the final analysis, the question of rationality must be settled by epistemological arguments *a priori*.²

Criteria for a Rational Reconstruction of Design Inferences

Integration

Rational reconstructions of design inferences should be integrated into a more general epistemic framework. This affects reconstruction at multiple levels. One's global theory of knowledge will place certain constraints on any sort of inference. Popperians, for example, will not tolerate inductive inferences of any sort; subjective Bayesians will reject the absolute eliminative inferences that characterize Popper's own approach; foundationalists will reject the psychological priors of the subjective Bayesians; and coherentists will fume at the linear nature of a foundationalist's inferential scheme. Conversely, if a strong case can be made for a particular reconstruction of design inferences, this will reflect unfavorably on theories of knowledge that render that reconstruction rationally impossible.

At a more detailed level, a successful reconstruction will locate design inferences among other inferences we make on a routine basis and will enable us to explain not only what makes them distinctive but also what fea-

² This is itself a methodological constraint that some philosophers, particularly those attracted by epistemic externalism, would reject. The defense of this position goes beyond the scope of this essay, but the idea that mere strength of subjective conviction is never a strong epistemic ground can be found across a wide cross section of epistemologists who share very little else in common, e.g., J. M. Keynes, *A Treatise on Probability* (London: Macmillan, 1963), 4; Karl Popper, *The Logic of Scientific Discovery* (New York: Harper, 1968), sec. 8; Carl Hempel, "Studies in the Logic of Confirmation," in *Aspects of Scientific Explanation* (New York: Free Press, 1965), particularly 8–10. The subjective conviction that subjective convictions are, in and of themselves, epistemically worthless is one of our most robust intuitions.

tures they share in common with other species of reasoning. The adequacy of a reconstruction may be judged in part from whether it treats intuitively similar types of inferences in a similar fashion.

Integration into a powerful and well-articulated framework can be an advantage for a rational reconstruction of design inferences. Conversely, any weaknesses of the framework will be a source of difficulties for the reconstruction as well. When we turn to an examination of the major contemporary models of design reasoning, we will see that this converse problem often determines the structure of the debate over a given model's adequacy.

Neutrality

Though the subject matter of our inferences (even of design inferences) varies widely, an adequate reconstruction will display the underlying structure of the inferences without reference to those topics. This is a fairly simple condition, but it can easily be misunderstood. Neutrality does not rule out the application of design inferences to specific topics—SETI, forensics, or even biology. But it does ensure that in the application the idiosyncratic features of a particular case will supply data that are relevant to the parameters of a topic-neutral formulation of the argument in question.

Neutrality can be satisfied by giving an elaborate formal account of the structure of design inferences, but by itself it does not require that the reconstruction be formal. An account of design inferences as explanatory inferences, for example, might appeal to criteria such as the simplicity of an explanation for which no adequate formal account currently exists. Keeping neutrality in mind may make us wary of nonformal elements in the reconstruction, but it need not make us paranoid.

One of the reasons that neutrality is important is that it safeguards us against the charge of special pleading either for or against particular conclusions. A reconstruction that is intrinsically gerrymandered, or permits gerrymandering, to sanction a particular bit of reasoning *because of its subject matter* will not and should not convince doubters. Conversely, a reconstruction that arbitrarily cuts out certain possible designers by invoking a distinction in one case that is selectively ignored in another will fail by the neutrality criterion as well.

Objectivity

Rational reconstruction would be a less urgent project if we never disagreed about the force of design arguments. One of our aims in reconstruction is to clarify, at a suitably generic level, the factors relevant to

the evaluation of design inferences so that in particular disputed cases we can make a dispassionate evaluation of their cogency. Different reconstructions provide us with differing sets of factors: some stress the simplicity or “loveliness” of an explanation, others the probability of the evidence given a hypothesis, and still others the degree of analogy between putative and known artifacts or the conformity of an event to an independently given pattern. But whichever reconstruction we opt for, the application of that account to particular cases will require us to make judgments about the relevance of the data to the parameters of that reconstruction. We will not have made serious progress if the application of the model to particular cases generates only irreconcilable clashes of intuition. A proper reconstruction should provide us with the resources to give an objective measure of the force of a design inference on which all informed and rational parties will agree. Disagreements regarding the cogency of a particular design argument should in principle be traceable either to differences in the relevant evidence available to the disagreeing parties or to specific inferential failures on the part of at least one of the disputants.

At first glance, the very project of rational reconstruction may seem to be at odds with the criterion of objectivity. The plausibility of a rational reconstruction of design inferences will depend in part on its ability to give us the “right” answers in cases where we have clear, shared intuitions. One might worry that by placing such weight on accommodating our instinctive judgments, the reconstruction will leave us with no impersonal standard in the evaluation of particular disputed design arguments. But this worry arises from a failure to separate the issues of verisimilitude and rationality. Our ability to account for our clear, shared intuitions counts in favor of the verisimilitude of the reconstruction. But the evaluation of specific arguments in light of that analysis is independent of our intuitive appraisals. The fact that there is widespread agreement regarding a large set of examples forms no part of the justification for the conclusion of a particular design inference. Thus, the tension between verisimilitude and rationality has a positive side: it is a guarantee that our project is not circular.

Some accounts of design inference promise a quantitative standard of comparison that would permit us to specify numerically, perhaps by probabilities, the degree of support a design hypothesis has on certain data. The potential for such a system, if we could have it, is great: all disputes about the relative strength of our inferences could in principle be resolved with mathematical precision. But there is also a cost in verisimilitude. The precision arrived at by a full mapping of degrees of support into the interval $[0, 1]$ is generally unrealistic as a reflection both of our intuitive judgments of force and of the precision and weight of our available evidence.

But the demand for objective measures of force need not impose a severe burden. At the simplest level the criterion of objectivity would be satisfied by a set of sufficient conditions for positive relevance—for instance, that evidence of a certain sort related to a design hypothesis in a certain way provides, *ceteris paribus*, evidence for that hypothesis. At a more sophisticated level, we might wish for a reconstruction that provides us with the resources to make rational comparisons: one that enables us to say, for example, whether the argument for D on the basis of *e* is stronger than the argument for D' on the basis of *e'* and, in particular, whether the argument for D on the basis of *e* is stronger than the argument for not-D.³ It does not strain credulity that a reasonable account of design inferences might provide us with at least a comparative measure of objective force.

Sensitivity

Design inferences are sensitive to various sorts of relevant information, and any adequate reconstruction will reflect this fact. One method of testing a reconstruction is to see whether it leaves out intuitively relevant parameters—whether, in other words, the objective evaluation according to the reconstruction is unchanged by the introduction of information that clearly ought to make some difference to our estimation of the argument's force.

Failures of sensitivity can take two forms. On the one hand, the reconstruction may simply fail to mention an epistemically relevant parameter. If the reconstruction itself does not claim to give an exhaustive list of parameters, then such a failure shows that the reconstruction is (at best) incomplete. On the other hand, the reconstruction may rule out the parameter altogether; and in that case it cannot be amended by any simple extension.

Our initial application of the sensitivity criterion is apt to arise from our intuition that a particular excluded parameter is clearly relevant. To the extent that our intuitive judgments can be relied upon, failure of sensitivity will suggest that the reconstruction in question offers a false picture of the inference. But the issues here are difficult: the tension between rationality and verisimilitude can pull us in either direction. An intrinsically compelling reconstruction might cause us to reconsider our intuitive judgments of relevance. The best critical use of the sensitivity criterion will rest not solely on intuitive judgments of relevance but also on the intrinsic merits of an alternative reconstruction that those judgments suggest.

³ Two classic discussions of qualitative, comparative and quantitative concepts of confirmation are Hempel, "Studies in the Logic of Confirmation," and Rudolf Carnap, *Logical Foundations of Probability*, 2nd ed. (Chicago: University of Chicago Press, 1962), 21–3 and 462ff.

Relevance

The criterion of relevance is the converse of sensitivity: an adequate reconstruction ought not to contain irrelevant parameters. We can determine relevance in two ways. First, we can vary the parameters of a particular inference and test our intuitive responses to them. Though this is subject to the fallibility of intuition—and we should be aware, in particular, that some intuitive judgments of probability are notoriously controversial—it can generate consensus when there is a large shared set of intuitions.

Second, if the reconstruction is intrinsically compelling, we can dispense with the examples. It happens frequently in mathematics that an example suggests a theorem that can later be demonstrated on wholly independent grounds, perhaps by mathematical induction. Similarly, a set of intuitively clear cases where a parameter is irrelevant may suggest that a correct reconstruction should leave that parameter out. But ideally the reconstruction will itself provide a compelling rationale for what it includes and excludes.

Accessibility

There is little merit in reconstructions that appeal to information patently inaccessible to us. We are, after all, interested in reconstructing *our* inferences, and within limits we know very well what we can and cannot do. Reconstructions that involve eons of laborious computation, for example, may show that a given inference is rational (in an ideal sense); but the availability of such computationally intensive machinery to an inferentially omniscient being says nothing about human rationality.

Accessibility is not only a problem with respect to computational resources. Bayesian reconstructions, for example, may appeal to prior probabilities. If objectively grounded priors are inaccessible, then such reconstructions could be caught on the horns of a dilemma, failing on either objectivity (if the Bayesianism is of the subjective variety) or accessibility (because objective priors that would do the epistemic heavy lifting are unavailable). We will return to this example later.

Power

There is a large area of overlap between the intuitions of any two normal people regarding the reasonableness of various inferences. If the overlap were always complete, we would have no urgent practical need for a rational reconstruction of our inferences, since disagreements would never

arise. But our agreement is never complete, and the inferences over which we contend are often some of the most interesting and pressing ones. And in some cases, we are nearly unanimous in our uncertainty—we find ourselves uncertain whether a given inference is strong enough to rely upon or not. Thus, we may wonder not only whether a particular argument is cogent but also why we find it so difficult to answer that question.

A rational reconstruction need not explain why we sometimes disagree or why we are sometimes uncertain. But a reconstruction that offered such an explanation would be particularly powerful, in the sense that it would provide not only guidance in resolving our disputes but also illumination regarding the causes of our perplexities. A compelling explanation for our disputes and our doubts would have to be taken as some evidence that the reconstruction affording that explanation is on the right track. In this sense, the explanatory power of a rational reconstruction is another criterion of its verisimilitude.

Three rational reconstructions feature prominently in contemporary discussions of design: Dembski's eliminative design inference, a reconstruction of design arguments as inferences to the best explanation, and a Bayesian reconstruction in which the concept of likelihood plays a prominent role. Traditionally, there has been a fourth construal of design as an argument from analogy, and considerations of analogy continue to play a role in some of the contemporary reformulations. But these modern reconstructions have certain virtues over the older reconstruction of design as a straight analogical argument.⁴ Our focus here is on the adequacy, in terms of the criteria proposed, of the three most significant current approaches to design reasoning.

Dembski's Eliminative Design Inference

The Basic Conception of the Explanatory Filter

The most detailed contemporary rational reconstruction of our design reasoning is the eliminative approach developed at length by William A. Dembski. In his book *The Design Inference*, Dembski represents the structure of design reasoning by an explanatory filter in which we ask whether the event in question is one of high, intermediate, or small probability.⁵ High probability events do not signify design: they are the province of law, the outcomes of natural regularity. Intermediate probability events may be ascribed to chance. Events of small probability—just how small is a matter Dembski treats separately—may still be ascribed to chance provided that

⁴ For a discussion of the analogical version of the argument and Hume's criticisms of it, see Elliott Sober, *Philosophy of Biology*, 2nd ed. (Boulder, CO: Westview, 2000), 33–5.

⁵ William A. Dembski, *The Design Inference* (Cambridge: Cambridge University Press, 1998).

they do not conform to an independently given pattern. But if they conform to such a pattern, they are said to be *specified*; and in that case, the chance hypothesis is eliminated. Specified small probability is, according to Dembski, a reliable marker of intelligent agency.

One of the ideas at the heart of this reconstruction is the probability of an outcome given the chance hypothesis. This parameter is intuitively relevant to any inference in which we are concerned to reject a chance hypothesis. The lower the probability of an outcome on chance, all else being equal, the more strongly we are inclined to suspect that something besides chance is in operation. If I flip a quarter twice and get two heads I am not inclined to take this as serious evidence that the quarter is biased in favor of heads; but if I get one hundred consecutive heads in as many tosses, my astonishment will be so great that I will likely be moved to doubt the fairness of the coin. Such an outcome is logically possible with a fair coin, but it is wildly improbable. Any reconstruction of design reasoning that fails to take this parameter into account will fail to meet the criterion of sensitivity.

Three aspects of Dembski's complex reconstruction deserve careful consideration. The first has to do with his notion of *specification* and some doubts about its epistemic relevance. The second arises from the purely eliminative nature of design inferences as he accounts for them and the charge that this leaves his account insensitive to some relevant parameters. The third is a cluster of concerns regarding Dembski's use of pragmatic considerations and the impact of these on the objectivity and neutrality of his reconstruction.

Specification and Relevance

The notion of specification bears a great deal of weight in Dembski's reconstruction. He explicitly relativizes it to the state of knowledge of given would-be design inferers, so that those with greater cognitive resources may be able to specify an event more easily than their less informed or less gifted counterparts. The basic criterion for separating "detachable" patterns, those that count as genuine specifications, from those that are not is this:

Given an event, would we be able to formulate a pattern describing it if we had no knowledge which event occurred? Here is the idea. An event has occurred. A pattern describing the event is given. The event is one from a range of possible events. If all we knew was the range of possible events without any specifics about which event actually occurred, could we still formulate the pattern describing the event? If so, the pattern is detachable from the event.⁶

⁶ Dembski, *The Design Inference*, 15. There is a more extended discussion in Dembski, *No Free Lunch: Why Specified Complexity Cannot Be Purchased without Intelligence* (Lanham, MD: Rowman and Littlefield, 2002), 15–18.

Dembski's illustration for this is the case of Nicholas Caputo, a clerk who had the job of selecting—ostensibly by a fair random process—the order in which candidates would appear on ballots in Essex County, New Jersey. It is well known that the candidate whose name appears at the top of the ballot has a small but statistically significant boost at the polls. Under the circumstances, it would raise suspicions of dishonesty if the proportions of Democrats and Republicans to get top billing were not roughly equal. And indeed they were not: over the course of several decades and forty-one elections in which Caputo was the sole official responsible, Democrats led the ballot forty times. For the sake of concreteness, let us suppose that the pattern looked like this:

ℳ: DDDDDDDDDDDDDDDDDDDDDDDDDDRDDDDDDDDDDDDDDDDDDDD

The pattern **ℳ** is specified, according to Dembski; it is a cheating pattern.

But given the definitions in *The Design Inference*, specification of a moderately low probability event is not difficult to achieve even when there is no obvious pattern of this sort. As Robin Collins points out in a review essay, a sufficiently powerful computer could generate a substantial proportion of the possible binary patterns corresponding to Caputo's possible assignments of Democrat or Republican to the top line of the ballot in successive elections.⁷ Those assignments would then be specified—with the awkward result that even a seemingly innocuous assignment such as

J: DRRDRRRDDDRDRDDDRRRDRRRDRRRDRRRDRDDDRDRDD

could now qualify as a specified event of small probability that would, in the absence of some further safeguard, suffice to convict a manifestly innocent election official of tampering.

One might suspect that mere conformity to an independently given pattern must not be all there is to specification—that the match between the computer's output and the hypothetical bland pattern of Democrats and Republicans given above is not what Dembski has in mind. But in his initial response to Collins he actually agrees that computer-generated patterns would be specifications. Since he also shares the intuition that bringing a computer into the courtroom would not suffice to convict someone who had produced the nearly even mix of Democrats and Republicans given above, Dembski balances the ease of specification by arguing that the probability required to justify a design inference is a function of the available specificational resources. The more resources available, the greater the chance for a false positive. In order to control false positives, therefore, we must

⁷ Robin Collins, "An Evaluation of William A. Dembski's *The Design Inference*: A Review Essay," *Christian Scholar's Review* 30 (2001): 329–41. Another penetrating, if unsympathetic, criticism of Dembski's original reconstruction from a broadly Bayesian standpoint is Fitelson, Stephens, and Sober, "How Not to Detect Design—Critical Notice: William A. Dembski, *The Design Inference*," *Philosophy of Science* 66 (1999): 472–88.

lower (that is, make stricter) the probability boundary on the event so that the resulting probability of our getting a false positive given the overall probabilistic resources is less than .5, the critical value for inferring design.⁸

There is obviously something fishy about Caputo's actual assignment \mathcal{C} . According to Dembski, if a computer were brought into the courtroom and generated a great number of patterns, including \mathcal{C} , then it would be too difficult to get a conviction since the required small probability boundary would be driven down in response to the inflation of specificational resources. Therefore, in his view, the court would be disinclined to allow a computer into the courtroom to generate the patterns. His comments here are worth quoting at some length.

Why, after all, should computers be allowed to inflate specificational resources in the Caputo case? As I pointed out in *TDI*, probabilistic resources are typically chosen according to pragmatic considerations that balance the need to eliminate chance when chance is not operating with the need to avoid eliminating chance when chance is operating. In civil trials, for instance, we allow far bigger probabilities to convict someone than in criminal trials (cf. "preponderance of evidence" vs. "to a moral certainty and beyond reasonable doubt") and thus require far fewer probabilistic resources than in criminal trials because mistakes in finding against someone are presumably less severe in civil trials than in criminal trials. And even in criminal trials one is not going to want to let probabilistic resources get too large for otherwise one will never convict anybody (the greater the number of probabilistic resources, the smaller the probability to eliminate chance—which here means conviction). Thus in bringing people like Caputo to justice for committing "probabilistic fraud," a court will be disinclined to increase probabilistic resources by using a computer to generate patterns.⁹

I find this analysis unpersuasive. According to Dembski, the judges would be disinclined to admit computer-generated patterns in the Caputo case because the patterns would be relevant and their presence would make it too hard to convict him. On the contrary, it seems obvious that the computer printout would have no relevance whatsoever to the question of Caputo's innocence or guilt. This is not a matter of variable pragmatic purposes; the existence of a specifying pattern *per se* does not make any difference to the strength of the case against Caputo. If this intuition can be

⁸ Dembski, *The Design Inference*, 184–98.

⁹ William A. Dembski, "Detecting Design by Eliminating Chance: A Response to Robin Collins," *Christian Scholar's Review* 30 (2001): 354.

sustained, then there is a serious problem with the relevance of Dembski's central notion of specification.¹⁰

In several subsequent publications, Dembski has modified his original account of specification to provide a different response to Collins's objection.¹¹ In *No Free Lunch*, he stresses that one's background knowledge must uniquely and univocally determine the set of outcomes that will serve to allow the rejection of the chance hypothesis, and he argues that a computer that generates all possible combinations of Ds and Rs in the Caputo example would be identifying not a unique rejection region but rather an ensemble of rejection regions.¹² So providing the court with the computer's output would not after all ruin the case against Caputo—a result that accords better with common sense.

But the new proposal is still vulnerable to a modification of Collins's original objection. Suppose that the election official is innocent and honest and that, in fact, the innocuous pattern *J* has occurred. Allow the computer to generate all of the possible patterns of Rs and Ds, including *J*, and then send out small packets of these patterns to individual people. Each recipient now has a rejection region ready to hand; he can reject chance if the actual order of Rs and Ds matches one in his packet. Someone will receive the packet containing, *inter alia*, *J*, and this person's background knowledge will enable him to name a rejection region (the set of patterns corresponding to the ones in his packet) that contains the actual outcome. But it is highly counterintuitive to say that he is in a position to declare that the innocuous pattern is evidence of cheating. On the other hand, reverting to the actual Caputo case with the highly dubious pattern *C*, it does not seem reasonable to declare a mistrial simply because all of these patterns have been sent out. The proper response is surely that all of this is quite beside the point. Here again, the relevance of the notion of specification seems dubious.

In *The Design Revolution*, Dembski makes a third attempt to block such counterexamples by admitting that, for design inferences, the combination of specification and low probability is insufficient: the specification must, in addition, be of "low complexity," a notion that he explicates in terms of algorithmic compressibility.¹³ There are certainly some cases where this constraint seems to be satisfied. Caputo's actual ballot line selections provide an example; the only way to get a simpler string of the same length as *C*

¹⁰ It is important to keep separate the use of the term "specification" as a place holder for whatever it takes, on analysis, to make an inference go through (*The Design Inference*, 5–6) and the actual definitions of the term given later by Dembski.

¹¹ I am greatly indebted to Lydia McGrew for help in the identification and analysis of the changes in Dembski's position discussed in the following paragraphs.

¹² William A. Dembski, *No Free Lunch*, 82–3.

¹³ William A. Dembski, *The Design Revolution: Answering the Toughest Questions about Intelligent Design* (Downers Grove, IL: InterVarsity, 2004), 83–84, 237.

would be for all forty-one selections to be identical. On the other hand, a random pattern of Ds and Rs such as \mathcal{J} is not compressible and therefore, presumably, not of low complexity. So with this restriction in place, the modification of Collins's example does not result in either a false negative or a false positive.

But there is more to be done in the justification of a constraint like this than to save a handful of intuitively correct inferences, and indeed this constraint seems *ad hoc* within the context of classical statistics. Moreover, this third proposal faces at least three difficulties. First, we need some fix on just how low the complexity of the specification ought to be. This is a more difficult matter than it might first appear. It will not do to say that the inference is stronger or weaker according as the specification is of lesser or greater complexity, since for Dembski the design inference is an all or nothing matter: the inference is *triggered* by the satisfaction of the relevant criteria. So we need an absolute cutoff. But Dembski does not give us a principled way of determining what that cutoff should be. If it turns out to be a function of our interests and needs, this will add another pragmatic layer to Dembski's reconstruction; and pragmatic elements, as we shall see below, tend to compromise the objectivity of the inference.

Second, the requirement seems far too strong. Forensic science abounds with incompressible patterns. Fingerprints, dental imprints and DNA samples are all highly complex, and yet they are often key pieces of evidence in criminal trials. We can tell instantly that novels and software code are the products of intelligent agency, though neither *War and Peace* nor Microsoft Word is algorithmically compressible. Personal passwords, social security numbers, and birth dates are all patterns that would apparently fail to satisfy the "low complexity" condition, but they are all readily accessible to the individuals who use them as the basis for numerous design inferences.

Finally, the requirement of low complexity fails to address cases in which a design inference is reasonable but in which the outcome can be described by two different patterns of roughly equal complexity, where one is clearly relevant and the other is clearly irrelevant. For example, suppose that Robert Mugabe can be described as a former student of one of my colleagues; he can also be described as the president of Zimbabwe. It should go without saying that only one of these facts is relevant to the inference that something nonrandom occurred when Mugabe himself won a lottery in 2000—run by a bank partially owned by the state. The difference, as far as an inference to shady dealings is concerned, is one of kind, not one of degree. Nothing about the relative complexity of the two descriptions is pertinent. The revised version of Dembski's system seems to be invoking a parameter that is, at least in some clear cases, irrelevant; and this leaves the relevance of low-complexity specification for design inferences in doubt.

Eliminative Inference and Sensitivity

By building his reconstruction on the classical concept of statistical inference, Dembski has created a model of design reasoning with a striking feature: it is purely eliminative. Chance is not *compared* to design to see which is the better hypothesis but is rather swept aside for its intrinsic failure as an explanation for specified events of small probability. Much of the controversy surrounding Dembski's approach comes down to this: is the fact that his reconstruction avoids treating design as a hypothesis in its own right a strength, since information about (say) the probability of the event on the hypothesis of design is arguably inaccessible, or is it a weakness since it makes his reconstruction insensitive to intuitively relevant parameters?

Dembski leaves no doubt that his account dispenses with some components widely thought to be crucial for probabilistic inferences. Contrasting his version of design reasoning with the Bayesian account, he makes a point of spelling out the ways in which his approach is more minimal.

Of the three types of probabilities that appear in Bayes's theorem—posterior probabilities, likelihoods, and prior probabilities—only one is relevant to the design inference, namely, the likelihoods. The probability of an event given the chance hypothesis is the only type of probability we need to consider. Posterior probabilities and prior probabilities play no role. This I take to be a huge advantage of the design inference. Posterior probabilities can typically be established only via prior probabilities, and prior probabilities are often impossible to justify.¹⁴

We can see more clearly what is at stake if we return to the Caputo case and attempt to give an account of it that makes use of the parameters Dembski wants to avoid. The likelihood on design is the probability of the expected outcome given deliberate intervention by the agent in question. How likely is it, then, that we would have gotten a severe imbalance in the awarding of lines to Democrats and to Republicans if indeed Caputo had cheated?

Dembski gives us a clue when he discusses the background information in terms of which, on his account, we formulate a detachable "cheating pattern" that specifies the actual event. The information includes:

- (1) Nicholas Caputo is a Democrat.
- (2) Nicholas Caputo would like to see Democrats appear first on the ballot since having the first place on the ballot line significantly boosts one's chances of winning an election.

¹⁴ Dembski, *The Design Inference*, 68; the same criticism appears in *No Free Lunch*, 102 and again in *The Design Revolution*, 240.

- (3) Nicholas Caputo, as election commissioner of Essex County, has full control over who appears first on the ballots in Essex County.
- (4) Election commissioners in the past have been guilty of all manner of fraud, including unfair assignments of ballot lines.
- (5) If Caputo were assigning ballot lines fairly, then both Democrats and Republicans should receive priority roughly the same number of times.¹⁵

It is not difficult to see how this information might make us suspicious regarding probabilistic irregularities in Caputo's assignment of ballot lines. The question, however, is whether the suspicion arises merely because we formulate from it a pattern that specifies the actual outcome or whether it rather functions to give us a relatively high probability for the slanted outcome on the hypothesis that Caputo circumvented the chance mechanism and took matters into his own hands. One argument in favor of the latter reading is that by amending the background information we appear to be able to alter the force of the inference. Suppose it were also part of our background knowledge that

- (6) Caputo suffers from a mental disorder that makes him mix up the names of current political candidates when he tries to say, read or write them, substituting Republicans for Democrats and *vice versa*.

Emending the background information by adding (6) would at least marginally weaken the reasonableness of the inference, leaving us less confident of Caputo's guilt and more puzzled. This cannot be explained if we use the background information only for generating a pattern, since the pattern can perfectly well be generated from (1)–(5) regardless of whether (6) is included. The natural way to account for the intuitive weakening of the argument is to see (6) as lowering the probability of a preponderance of Democrats over Republicans on the assumption that Caputo had cheated. A high likelihood of a disproportionate outcome given that Caputo was trying to cheat seems crucial to the force of the inference.

A similar argument can be made for the relevance of prior probabilities. Suppose that instead of adding (6), we modified the background information by adding:

- (7) Caputo has repeatedly shown himself to be a pietist of stern and upright moral character who deplores dishonesty and nepotism.

This does not affect the likelihoods: were Caputo to break character and try to cheat, we would expect something more or less like what he actually did. It does, however, lower the prior probability that Caputo would cheat at all. We might be inclined to think him guilty in any event: even people of known probity may break character through weakness of will from time to time.

¹⁵ Dembski, *The Design Inference*, 16.

But the stronger our evidence for (7), all else being equal, the less confident we ought to be in inferring Caputo's guilt on the basis of the actual outcomes. We could depress the prior probability still further by adding items such as

- (8) Secret FBI video tapes of Caputo selecting candidates on multiple occasions indicate strongly that he actually used the urn model fairly and expressed visible annoyance when it kept producing Democrats at a higher than expected rate.

It appears that we could, by the introduction of sufficient background information relevant to the probability of Caputo's cheating, depress the prior probability so far that it might in the end be more reasonable to suspect that Caputo had experienced a long run of unusual luck than to think that he had cheated.

With respect to both the likelihood on design and the prior probability of cheating, then, there is reason to believe that Dembski's eliminative design inference is insensitive to relevant information—not insensitive by amendable omission, but insensitive in the strong sense that Dembski has deliberately blocked any appeal to such parameters. When we turn to an examination of Bayesian reconstructions of the design inference, we will have to consider his countercharge that these parameters are typically inaccessible.

Pragmatic Considerations, Objectivity, and Neutrality

In many scientific applications of statistical reasoning, our desire to control for false positives is conditioned by factors of a pragmatic nature. The head of a research team that has invested heavily in the development of a new vaccine may have evidence that he runs only a 10 percent risk of being proved wrong if, on the basis of the trials run so far, he declares it to be effective. But if the company stands to lose an enormous amount of money in future contracts for making an embarrassing error, then there may be a very good practical reason to withhold any announcement pending further clinical trials.

In applying the statistical model to design inferences, Dembski also stresses the importance of pragmatic considerations. The determination of the relevant set of probabilistic resources in terms of which the small probability boundary must be set will depend, he argues, on how important it is for us to avoid false positives; and individuals with different purposes or standards will set their boundaries differently.¹⁶ In the passage we have

¹⁶ *Ibid.*, 203–7; cf. *No Free Lunch*, 49, 62, 72.

already quoted he invokes such considerations to explain why we require a higher standard of evidence in a criminal trial than in a civil trial.

Such considerations must, of course, be taken into account in guiding our actions. But it is not apparent why they should be thought to be of any *epistemic* relevance. Suppose, to elaborate on Dembski's example, that I am a juror in a criminal trial where the preponderance of evidence points to the defendant's guilt, but the evidence is not quite so strong as to place the matter beyond reasonable doubt. Wise provisions in law mandate that I not vote for conviction without a more powerful case, since it is far worse to condemn an innocent man on criminal charges than to let a guilty one go free. But all of this touches only on actions. May I not still believe, and believe for good reason, that the accused was guilty? Surely so: nor am I doing him any injustice in holding the belief. Since the evidence indicates his guilt, it would not be unreasonable for me to believe it.

The difficulty here is that according to Dembski there might be two or more individuals with access to the same relevant data only one of whom could infer design from those data, since the other has a stronger motivation to avoid false positives. Pragmatic issues create a sliding scale of probabilities on chance, and the data will trigger or fail to trigger design inferences depending on the individual's goals and purposes. But this runs afoul of the objectivity constraint on rational reconstructions. Disagreements over the inferrability of design in a given case must be traceable to epistemically significant factors such as differences in relevant evidence or inferential slips on the part of one party or the other. Different pragmatic goals and purposes do not carry the requisite epistemic weight.

If there were an objective way to determine which set of probabilistic resources is relevant to the occurrence of an event, we might salvage objectivity by requiring all reasoners to take precisely that set of resources into account. But in Dembski's system there is no *rationaly* privileged set: resources are relevant to different reasoners according to their "interests and needs."¹⁷ This is no mere omission on Dembski's part: he specifically criticizes Borel for having selected a probability boundary in a way that "neglects an inquirer's interests and context."¹⁸ In light of this heavy reliance on the variability of interests, needs and contexts, there appears to be no objective way to determine whether p_{Ω} —the probability of a false positive, taking all relevant probabilistic resources into account—is greater than 1/2 or not.¹⁹

¹⁷ Dembski, *The Design Inference*, 191.

¹⁸ *Ibid.*, 6–7. The discussion refers to Emile Borel's 1943 discussion of probability boundaries, available in an English translation by Maurice Baudin as *Probabilities and Life* (New York: Dover, 1962). Dembski is criticizing pages 25–32 of Borel's work, to which nevertheless his own approach evidently owes much.

¹⁹ Dembski, *The Design Inference*, 194.

Can we plug this hole by appealing to a maximum set of probabilistic resources? Dembski does provide an argument to the effect that the greatest possible number of specifications (given some physical assumptions) is bounded by 10^{150} . This represents, as he sees it, an upper bound on the number of physically possible specifications in the lifetime of the universe. If we take all of these resources into account and employ a probability boundary of $1/2 \times 10^{-150}$, he suggests, then design inferences that meet this universal probability boundary are always foolproof: there is nothing further to be factored in that might reduce the cogency of the inference.²⁰

Unfortunately, the universal probability bound does not address the problem of objectivity, for two reasons. First, it is defined explicitly in terms of possible specifications, and as we have already seen there are serious reasons to doubt that specification in and of itself is of any epistemic relevance. But a second problem with an appeal to this boundary is that even Dembski's own examples of design reasoning do not meet it. The probability of getting the pattern \mathcal{C} using a fair chance process is low, but it is nowhere close to the universal boundary. Something similar goes for many of his other paradigmatic design inferences. The evidence for design in each case seems very strong, even rationally compelling; but since the probabilities fall drastically short of the only boundary Dembski discusses that is not sensitive to pragmatic and personal considerations of context and need, there are no resources for explaining why they could not be rationalized away by an appropriate choice of probabilistic resources. Thus, even the invocation of the universal probability bound does not rescue Dembski's reconstruction from a collision with the criterion of objectivity.

The heavy dependence on pragmatic considerations also raises questions concerning the neutrality of Dembski's reconstruction. This is particularly evident in his discussion of the rationale for excluding a computer from the courtroom in the Caputo case. Even granting for the sake of argument that the inflation of probabilistic resources would have an impact on the reasonableness of a judgment regarding Caputo, Dembski's invocation of the desire to get at least some convictions seems to be a classic case of gerrymandering the inference because of its subject matter so as to make it easier to infer design.

Classical Statistics and Integration

Dembski's reconstruction of design inferences is very clearly an extension of the classical statistical procedure of rejecting a null hypothesis in

²⁰ Ibid., 212. See also Dembski's comments in "Detecting Design by Eliminating Chance" 355–6, and *No Free Lunch*, 83.

terms of an independently specified rejection region in the set of possible outcomes. One might suppose that this would provide at least a *prima facie* integration for his eliminative account, so that anyone wishing to level broad-scale criticisms against his entire approach would have to take on not just Dembski but the entire school of classical error statisticians.

In a sense, it does. But at the same time this tight linkage raises some questions of its own. Classical statistics neither provides nor pretends to provide a global theory of rational inference. Jerzy Neyman, one of its most notable exponents, forswore the attempt to draw inferences and viewed significance testing as a means rather to optimizing behavior.²¹ To the end of his life R. A. Fisher resisted this pragmatizing of statistical inference²²; and perhaps Dembski, notwithstanding his frequent appeals to pragmatism, would wish to join him. But even Fisher conceded that a tremendously improbable result might have to be conceded to have been due to chance if the prior probability of chance were great enough.²³

Among the classicists with an optimistic outlook about scientific inference, significance testing is not uncommonly viewed as providing only one piece of the inferential puzzle. W. S. Gosset, whose *t*-test features so prominently in Fisher's writings, writes that a sample far from the mean of the ostensible population distribution

doesn't in itself necessarily prove that the sample is not drawn randomly from the population even if the chance is very small, say .00001: what it does is to show that if there is any alternative hypothesis which will explain the occurrence of the sample with a more reasonable probability, say .05 (such as that it belongs to a different population or that the sample wasn't random or whatever will do the trick) you will be very much more inclined to consider that the original hypothesis is not true.²⁴

This frank admission of the need to consider the likelihoods on alternative hypotheses is fundamentally at odds with the method Dembski has developed.

The reason for this internal strife among the classicists is that significance testing, considered in and of itself, appears to be insensitive to information we need in order to arrive at a rational judgment of the credibility of our hypotheses. We may, like Neyman, despair of acquiring the required

²¹ This comes out with particular clarity in Neyman's papers "'Inductive Behavior' as a Basic Concept of Philosophy of Science," *Revue de l'institute internationale de statistique* 25 (1957): 7–22, and "Frequentist Probability and Frequentist Statistics," *Synthese* 36 (1977): 79–131.

²² See Fisher, *Statistical Methods and Scientific Inference*, 3rd ed. (New York: Macmillan, 1973), 103–7.

²³ *Ibid.*, 42–4.

²⁴ From a letter to Egon S. Pearson, dated May 11, 1926, printed in Pearson's appreciation of Gosset, "'Student' as Statistician," *Biometrika* 30 (1939): 243.

additional information and give up all hope of constructing a theory of inductive inference. But if we aspire to something more, then like Gosset—and Fisher himself—we will have to seek for it outside the narrow boundaries of significance testing.

Design Reasoning as Inference to the Best Explanation

The idea that design reasoning might be construed as an explanatory inference has deep historical roots. Although the explicit discussion of such reasoning was catalyzed by Charles Sanders Peirce's investigations in the latter half of the nineteenth century, there are intimations of its applicability to design arguments nearly a century earlier in the classic work of William Paley. Recent work that reconstructs design arguments as inferences to the best explanation has not yet achieved the level of detail exhibited by Dembski's system. But the idea can be found in the writings of several contemporary philosophers of science, including both advocates²⁵ and critics²⁶ of the argument from design.

The Basic Conception of Inference to the Best Explanation

In a classic formulation, Peirce suggested that inference to the best explanation is initiated by a mental act of "abduction"²⁷ that conforms to the following schema:

The surprising fact C is observed;
But if A were true, C would be a matter of course,
Hence, there is reason to suspect that A is true.²⁸

Here C is the unexpected bit of evidence that initiates the inference, and A is an explanatory hypothesis that would, if true, remove our surprise at C. In numerous places Peirce calls this an "inferential step," although the con-

²⁵ See Stephen C. Meyer, "The Methodological Equivalence of Design and Descent," in *The Creation Hypothesis*, ed. J. P. Moreland (Downers Grove, IL: InterVarsity, 1994), particularly 88–98, as well as Moreland's own introduction to the volume, particularly 26–7.

²⁶ Sober, *Philosophy of Biology*, 30–3.

²⁷ At various times Peirce refers to this method of reasoning as "retroduction," "hypothesis," "presumption," and "originary argument."

²⁸ C. S. Peirce, *Collected Papers*, ed. C. Hartshorne and P. Weiss (Cambridge, MA: Harvard University Press, 1935), 5.189.

clusion—that there is reason to suspect that A is true—is a good deal weaker than the outright conclusion that A.²⁹ As he says elsewhere, abduction merely suggests that something *may be true*.³⁰

Abduction alone cannot do more than generate hypotheses that cover the facts in question. But obviously many hypotheses may do this much, and as these may be mutually incompatible they cannot all be true. Recent work on inference to the best explanation, or IBE, has stressed the importance of a second stage in which we select from among the competing hypotheses by means of criteria that reflect explanatory virtues.³¹ Paul Thagard suggests that at least three criteria are crucial: consilience (the capacity to explain diverse independent classes of facts), simplicity (the capacity to explain the facts at hand without invoking a host of auxiliary hypotheses of narrow application), and analogy (the invocation of causes of a type known to have explanatory value in similar contexts).³² Thagard argues that these criteria conform to actual scientific practice and shed more light on scientific inference than the old hypothetico-deductive model of scientific inference could afford.

In his book *Inference to the Best Explanation*, Peter Lipton offers a similar list.³³ “Lovely” theories, according to Lipton, stand out from their rivals because of the detail with which they articulate the causal mechanism invoked,³⁴ the level of precision and detail with which they explain the data,³⁵ their parsimony in invoking fundamentally new types of phenomena,³⁶ and their intrinsic elegance and simplicity in an unabashedly aesthetic sense of those terms.³⁷

²⁹ There is a useful overview of Peirce’s development of abduction in Ilkka Niiniluoto, “Defending Abduction,” *Philosophy of Science* 66 (Proceedings) (1999): S436–S451.

³⁰ Peirce, *Collected Papers*, 5.171.

³¹ N. R. Hanson discusses the two stages in *Patterns of Discovery* (Cambridge: Cambridge University Press, 1958), 85–92, and develops the distinction more fully in “Is There a Logic of Scientific Discovery?” in *Current Issues in the Philosophy of Science*, ed. Feigl and Maxwell (New York: Holt, Rinehart, and Winston, 1961), 22–35. Gilbert Harman raises the problem of spelling out and justifying the criteria for selecting from among rival explanations in his essay “The Inference to the Best Explanation,” *Philosophical Review* 74 (1965): 88–95, but he does not make serious progress toward articulating the relevant criteria. The best recent work on IBE consists in attempts to fill this gap and subsequent criticisms and defenses of those attempts.

³² Paul Thagard, “Inference to the Best Explanation: Criteria for Theory Choice,” *The Journal of Philosophy* 75 (1978): 76–92.

³³ Peter Lipton, *Inference to the Best Explanation* (New York: Routledge, 1991), particularly 117ff.

³⁴ *Ibid.*, 118.

³⁵ *Ibid.*

³⁶ *Ibid.*, 119. See also Michael Friedman, “Explanation and Scientific Understanding,” *Journal of Philosophy* 71 (1974): 1–19, to whose discussion Lipton adverts.

³⁷ Lipton, *Inference to the Best Explanation*, 68.

Like Thagard, Lipton urges that we can make better sense out of scientific practice by conceiving of it in terms of IBE than we could otherwise. But Lipton also suggests that there is a significant connection between our judgments of the loveliness and the likeliness of a theory. The criteria of loveliness, in his view, are the features an argument has that lead us to say that the premises make the conclusion likely.³⁸ In sifting through potential explanations with the help of such criteria, we need not fear that we are abandoning the search for the likeliest explanation: likeliness is determined, at least in part, by just such explanatory considerations.³⁹

The dependence of likeliness on loveliness rather than the other way around is an important feature of Lipton's position, and it is unfortunate that he sometimes speaks of it in psychological and sometimes in logical terms. Psychologically, we may indeed be inclined to select explanations in accordance with something like Lipton's criteria—explanations that would, if true, afford us the most understanding—when our goal is to find the most likely explanation. This seems to be what Lipton has in mind when he suggests that we

want a model of inductive inference to describe what principles we use to judge one inference more likely than another, . . . we want our account of inference to give the *symptoms* of likeliness, the features an argument has that lead us to say that the premises make the conclusion likely.⁴⁰

But in the same context he expresses a more significant analytical concern: that the explication of loveliness in terms of likeliness would “push Inference to the Best Explanation toward triviality.” A deep version of IBE, he suggests, will show how explanatory considerations determine plausibility, and he insists that explanatory loveliness is “conceptually independent of likeliness.”⁴¹ A similar insistence on the priority of explanatory considerations turns up elsewhere in the IBE literature. John Foster, for example, makes it a key component of his solution to the problem of induction that “[t]he only primitive rational form of empirical (non-deductive) inference is inference to the best explanation.”⁴² Whether the strong asymmetry between likeliness and loveliness that makes IBE a fundamental rule of inference is a strength or a weakness of this approach is a serious question to which we will shortly return.

³⁸ *Ibid.*, 62.

³⁹ *Ibid.*, 63.

⁴⁰ *Ibid.*, 62.

⁴¹ *Ibid.*, 64.

⁴² John Foster, *A. J. Ayer* (New York: Routledge and Kegan Paul, 1985), 227. Compare Laurence Bonjour, *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998), 207ff.

Embedding Design Reasoning in IBE

A rational reconstruction of design reasoning as a species of IBE will have some distinctive characteristics. First, unlike Dembski's filter, IBE is explicitly contrastive. The explanatory virtues come in degrees, and rival hypotheses will have or lack these virtues to varying extents. Thagard lays great stress on the idea that IBE is a comparative mode of inference:

Inference to a scientific theory is not only a matter of the relation of the theory to the evidence, but must also take into account the relation of competing theories to the evidence. Inference is a matter of choosing among alternative theories, and we choose according to which one provides the best explanation.⁴³

There can be no question, then, of sweeping the field clear of rival theories and concluding design without examining the explanatory virtues of the design hypothesis itself. The best explanation—the one affording us “the greatest potential understanding,” to use Lipton's phrase—will be the one that exhibits a better constellation of theoretical virtues than its rivals.

Second, design reasoning on this model will involve considerations of causal adequacy.⁴⁴ Again, this contrasts sharply with Dembski's filter, where design is simply the set-theoretic complement of the disjunction of regularity and chance.⁴⁵ Asked to attribute the carving of the *Winged Victory of Samothrace* either to a desert sandstorm or to a newborn infant, we would doubtless suppose that some third alternative had been overlooked. But if (given some bizarre constraints on our abductive generation of hypotheses) those were our only live options, IBE would not endorse the infant hypothesis *simply* because a sandstorm has little chance of producing such a masterpiece.

How do intuitively strong design arguments fare when evaluated by the criteria for IBE? Taking the Caputo case as a test instance, we can make the following plausible assessments. For the initial abductive step, the design hypothesis (that the assignments are the result of deliberate cheating by

⁴³ Thagard, “Inference to the Best Explanation,” 91.

⁴⁴ Meyer stresses this in “The Methodological Equivalence of Design and Descent,” 90, 94. See also Meyer's essay “Evidence for Design in Physics and Biology: From the Origin of the Universe to the Origin of Life,” in *Science and Evidence for Design in the Universe*, ed. Behe, Dembski, and Meyer (San Francisco: Ignatius, 2000), particularly 94–6.

⁴⁵ This is not to say that Dembski construes specific design hypotheses noncausally: he maintains that design, as a category, sharply constrains our causal accounts of phenomena. But the inference to design in the first place, as opposed to the inference from the conclusion that something is designed to the existence of an intelligent agent causally competent to bring it about, is wholly free of causal considerations—deliberately so, to avoid “prejudicing the causal stories we associate with design inferences” (*The Design Inference*, 36).

Caputo) passes with flying colors. The ballot-line assignments are surprising, and they would be a “matter of course” if Caputo had cheated. Moving to the selection step, the cheating hypothesis is consilient insofar as it explains the outcomes of forty-one separate and (on the chance hypothesis) independent assignments of ballot lines. These all belong to a single type, it is true; but these are the data at hand, and no hypothesis under serious consideration explains a wider range of relevant facts. Its simplicity is admirable, its precision nearly as high as one could hope for, and its mechanism well understood. The invocation of an unscrupulous individual to commit an unethical act does nothing to swell our ontology, and the analogy with successful explanations in other cases is patent.

As far as rivals are concerned both chance and alternative design hypotheses make a poor showing, though they fail for different reasons. No known unintelligent process biases the drawings from a normal urn to that extent with any appreciable frequency: the chance hypothesis does not make it past the first abductive screening, for it fails to remove our surprise. Our background knowledge as given by Dembski mitigates against outrageous hypotheses involving tampering by other agents such as Democrat-loving aliens, since invoking them (and the auxiliary hypotheses necessary to account for the fact that we missed their repeated visits to New Jersey) suffers from a conspicuous failure of simplicity. In short, the cheating hypothesis is so clearly superior to its rivals that it is extremely difficult to resist the conclusion that it is true.

We can go further. Recall the various extensions of our background knowledge to which Dembski’s explanatory filter seemed insensitive. If Caputo were mentally deranged when it came to distinguishing the names of Republicans from those of Democrats, then the hypothesis of his trying to cheat for his own party’s benefit would no longer render the observed outcome a matter of course. If he were a stern pietist, the cheating hypothesis would suffer strain at the points of analogy and ontology; for we rarely find it helpful to invoke the existence of the morally upright as a causal explanation for venal acts. If our secret surveillance tapes of Caputo strongly suggested that he used the urn fairly, the cheating hypothesis would suffer at the point of mechanism. In every case, the criteria for IBE seem to track our intuitive judgments about Caputo perfectly.

Reconstructing design reasoning in the framework of IBE has evident attractions. It allows us to see design arguments in relation to a wide range of other inferences in the broad class of explanatory arguments, a class that, according to the advocates of IBE, comprises most scientific reasoning. They fall into a subclass of those arguments that appeal to causal hypotheses, and they are distinguished from other causal IBEs by their appeal to a particular sort of cause. Intuitively compelling inferences to

design are plentiful and, as the eternal popularity of Sherlock Holmes assures us, the reconstruction of such inferences in terms of “reasoning backward, or analytically” is particularly compelling.⁴⁶ One cannot ask for much more on the count of integrating a rational reconstruction with a broader theory of inference.

The criteria for the evaluation of design explanations will be the same as for any other explanation; the theoretical virtues carry positive epistemic weight regardless of subject matter, though the sort of evidence relevant to our judgments of an explanation’s theoretical virtues will be constrained by the nature of the hypothesis. And as we saw in our examination of the Caputo case, design arguments, when cast in IBE form, are sensitive to evidence regarding rival hypotheses in a way that a purely eliminative reconstruction is not. Thus, a reconstruction of design reasoning as an IBE can also number neutrality and considerable sensitivity among its merits.

There is, however, a price exacted for these virtues. Critics of IBE have leveled serious charges against it, two of which deserve particular notice. First, they charge that the initial abductive screening may provide us with such a paucity of hypotheses that our best available explanation may be only the best of a bad lot, so that the mere fact that we have discovered the best explanation from among those available to us is no indication that we have discovered the truth. If this criticism can be sustained, design inferences built on this model will be insensitive to at least some of the factors that separate viable from nonviable hypotheses.

Second, even waiving this, critics of IBE argue that the criteria proposed by its defenders have—our intuitions to the contrary notwithstanding—no intrinsic connection to truth. Aesthetic virtues may attract our attention and even our assent without genuinely confirming the hypotheses they adorn. If IBE is vulnerable to this charge, an explanatory reconstruction of design inferences will be struck down on grounds of irrelevance.

Abductive Screening and Sensitivity

Bas van Fraassen, one of the most outspoken and committed critics of IBE, gives a canonical formulation of the “best of a bad lot” argument in *Laws and Symmetry*:

Inference to the Best Explanation is not what it pretends to be, if it pretends to fulfill the ideal of induction. As such its purport is to be

⁴⁶ Sir Arthur Conan Doyle, *A Study in Scarlet*, Part 2, chap. 7. The methodology of Sherlock Holmes is examined in some detail in Umberto Eco and Thomas Sebeok, eds., *The Sign of Three: Dupin, Holmes, Peirce* (Bloomington, IN: Indiana University Press, 1983), particularly chaps. 2, 3, 9 and 10.

a rule to form warranted new beliefs on the basis of the evidence, the evidence alone, in a purely objective manner. It purports to do this on the basis of an evaluation of hypotheses with respect to how well they explain the evidence, where explanation again is an objective relation between hypotheses and evidence alone.

It cannot be *that* for it is a rule that only selects the best among the historically given hypotheses. We can watch no contest of the theories we have so painfully struggled to formulate, with those no one has proposed. So our selection may well be the best of a bad lot.⁴⁷

The burden of this complaint is that there is no guarantee that the initial abductive screening of hypotheses is more likely than not to include the true hypothesis among the live options. Even granting that the criteria of selection from this point onward are relevant and their application objective—granting, in fact, that the true explanation will make a better showing than all of the other possible rivals—it does not follow that the very best explanation is likely to be among those we have under serious consideration. We can examine only a tiny handful of the infinitely many hypotheses that might be proposed. Unless there is some reason to believe that the true hypothesis lies among that handful, IBE never gets past the first and fundamental hurdle: beliefs formed in accordance with it are not rational, since even if we are reasoning from known premises our conclusions are not likely to be true.

It has been urged several times that this type of objection would also undermine van Fraassen's own constructive empiricism.⁴⁸ For van Fraassen himself allows inference to discriminate among rivals so long as we refrain from making bold claims about the truth of the "best" theory and conclude only that it is the one most likely to be, in his terminology, empirically adequate. But this *tu quoque* does not really get to the heart of the problem. The most such a response can show is that the assumption required for a vindication of IBE—that there is *something* epistemically privileged about the set of beliefs we take to be serious contenders—is one van Fraassen also shares. But this does nothing to show that either side is right to make the assumption, nor does it show that van Fraassen's realist critics have the resources to sustain their strong claims about the extent of that privilege.

Peirce's formulation of the schema for abductive inferences gives us little comfort here. It is easy enough to generate scenarios in which the explanations that would do the most to remove our surprise are actually unreasonable to invoke, whereas the true explanation, though it does little to remove our surprise, is clearly the most reasonable one to accept. Suppose

⁴⁷ Bas van Fraassen, *Laws and Symmetry* (Oxford: Clarendon, 1989), 142–3.

⁴⁸ Lipton, *Inference to the Best Explanation*, 176. See also section III of Stathis Psillos, "On van Fraassen's Critique of Abductive Reasoning," *Philosophical Quarterly* 46 (1996): 31–47.

I casually pull a quarter out of my pocket and flip it five times, getting heads all five times—an outcome with a probability of about 3 percent on the hypothesis that the quarter is fair. If the quarter were two-headed, the result would be a matter of course. Yet in my experience two-headed quarters are tremendously rare. It would be more reasonable to assume that a coincidence had occurred than to invoke an improbable hypothesis in order to account for the mildly surprising outcome.

Something similar goes for conspiracy theories. They often tie together a certain set of details in a very neat package, and on that basis they may make a strong showing at the first stage of IBE. If we restrict our focus to the data that raise our eyebrows, then the postulation of a global Masonic cabal might well let us replace our initial surprise with the sense that this was only to be expected. But at the same time most conspiracy theories are so outlandish that we judge them to be nonstarters.

The quarter-flipping and conspiracy theory cases both suggest that the initial abductive screening, at least as Peirce describes it, fails to forge the right sort of link between our hypotheses and the truth. In consequence, a rational reconstruction of design reasoning as IBE looks vulnerable to the charge of insensitivity. But this very failure gives us a clue as to how IBE needs to be repaired in order to meet the best of a bad lot criticism. Removal of our surprise regarding the limited data that catch our attention is not sufficient: the question of which hypotheses to take seriously needs to be addressed in the light of our total evidence and not simply in terms of their ability to account for the surprising event at hand. Loveliness and likeliness are not identical concepts, as Lipton himself admits; an explanation may suffer a decrease in its likeliness when a plausible rival is introduced, but it will not typically lose any of its loveliness thereby.⁴⁹ When we want to know what to infer, we need to know what is likely, not merely what is lovely. And at the abductive stage this means that we must avoid tunnel vision in the generation of our hypotheses.

This seems reasonable as a piece of negative advice, but when we try to turn it into a positive program it raises a new question: by what standards should the potential explanatory hypotheses themselves be generated? It will not do to expand the scope of IBE to cover this gap, urging that the hypothesis to be selected must be the best explanation for every piece of our relevant background knowledge; that would stretch even the flexible concept of explanation past its breaking point.⁵⁰ The judgment that *it is not reasonable to infer that the quarter is two headed* is certainly sensitive to our

⁴⁹ Lipton, *Inference to the Best Explanation*, 62.

⁵⁰ Which is not to say that it has not been tried. See Gilbert Harman, *Thought* (Princeton: Princeton University Press, 1973) for an attempt to reduce literally all inference to IBE.

evidence, but it does not in any obvious way depend on the *explanatory* virtues or vices of that hypothesis and its rivals.

When we evaluated the Caputo case we suggested that we could accommodate information negatively relevant to the hypothesis of cheating by pointing out that the hypothesis would now suffer strain at the points of analogy and ontology. But in light of the “best of a bad lot” criticism, it appears that matters are not so easily dealt with. For if the initial abductive screening limits the set of hypotheses from which we are permitted to choose, then the most that we can achieve by balancing initial abductive loveliness with these negative factors is to reduce the overall attraction of the cheating theory; and this does not admit into the set of live options any new hypotheses that might, on balance, fare better.

The moral we should draw from this criticism and reply is that IBE, whatever its virtues and however wide its range of application, is not a fundamental rule of inference.⁵¹ The initial abductive step needs to be understood in accordance with a principle of total evidence rather than in terms of Peirce’s simple rule. And assessing the initial credibility of hypotheses will require something like a probabilistic framework for coordinating the various types of relevant information.

Theoretical Virtues and Relevance

The second primary line of attack on IBE grants for the sake of the argument that the initial abductive step generates a field of hypotheses with a reasonable chance of containing the truth and focuses instead on the theoretical virtues by which we select among them. We have already seen that loveliness and likeliness are not the same thing, and conspiracy theory cases indicate that they can come apart. Why, then, should the explanatory virtues be confirmatory?

It is not in dispute that reasoners are often swayed by the sorts of considerations outlined in our sketch of IBE: Thagard persuasively argues that consilience, simplicity, and analogy were guiding methodological principles for Lavoisier, Fresnel, and Darwin.⁵² But by itself this historical argument is ineffective for three reasons. First, at a deep level we cannot take it for granted that the theories thus arrived at are true. Thagard’s examples themselves underscore this. No contemporary chemist would be content with Lavoisier’s understanding of chemical bonding; Fresnel’s theory of the

⁵¹ For a similar conclusion arrived at by a somewhat different route, see Timothy Day and Harold Kincaid, “Putting Inference to the Best Explanation in its Place,” *Synthese* 98 (1994): 271–95.

⁵² Thagard, “Inference to the Best Explanation,” 80–1, 86–9.

ether, even enhanced by the Lorentz-Fitzgerald contraction, went out of vogue before the First World War; and though Darwin's own arguments for his position are of more than historical interest, they have taken a back seat to more sophisticated arguments for a more sophisticated version of the natural selection hypothesis of which Darwin never dreamed.⁵³ Second, even if we were convinced that many theories reached by attentive cultivation of the theoretical virtues are true, we would still need some method for deciding among theories that exhibit differing virtues and vices—more simplicity, say, but less consilience. Third, even if we had a principled way to evaluate such tradeoffs, it would not follow that the virtues are confirmatory by themselves. The relevance of a given theoretical virtue and its weight vis-à-vis another may be context dependent; and if so, we need a better understanding of how our background knowledge should be factored in than IBE, at least in the form we have considered so far, provides.

Keats notwithstanding, beauty is not truth; and even the staunchest defenders of IBE do not claim that it is. But if IBE is to stand on its own as a methodology, then we need an account of the link between lovely and likely theories that does more than appeal to the practice of scientists or the aesthetic attractions of notions like simplicity. It seems evident even from the conspiracy theory case that likeliness cannot be wholly a matter of loveliness, and Lipton himself retreats to the more modest position that likeliness can only be *partially* accounted for in explanatory terms.⁵⁴ There is a great deal of flexibility in a position thus qualified. It buys proponents of IBE an indefinitely large logical space into which they can retreat in the face of apparent counterexamples. But for just that reason it risks evacuating the program of all interesting content.

Recent criticisms suggest that a further retrenchment is necessary. For even when our judgments of likeliness and loveliness coincide, it is by no means clear that the likeliness is due to the explanatory virtues in question.⁵⁵ Consider precision, which Lipton includes among the explanatory virtues because, as he puts it, “we understand more when we can explain the quantitative features of a phenomenon, and not just its qualitative ones.”⁵⁶ The motivation for some such criterion is pretty clear. Predictions of soothsayers and astrologers are notoriously vague: almost everyone can find something corresponding to his life under any sign in the newspaper's weekly horoscope. By taking precision into account, we can see that the “accuracy” of these predictions is purchased at the cost of a desperate loss of precision,

⁵³ See John Maynard Smith, *The Theory of Evolution*, 3rd ed. (Cambridge: Cambridge University Press, 1993).

⁵⁴ Lipton, *Inference to the Best Explanation*, 64.

⁵⁵ Eric Barnes argues for this in some detail in “Inference to the Loveliest Explanation,” *Synthese* 103 (1995): 251–77.

⁵⁶ Lipton, *Inference to the Best Explanation*, 118.

so that the truth of astrology does not count as a lovely explanation for, say, my finding that “the week ahead holds many changes in store.” On the other hand, the verification of a highly precise prediction—“A man wearing a black beret will sit down on the park bench at 11:47 and depart five minutes later, leaving behind a brown paper bag containing \$40,000 in unmarked bills”—would seem to be very well explained by the hypothesis that my informant knows what he is talking about.

But as Eric Barnes points out, the confirmational benefits of precise predictions are already underwritten by some basic results in probability theory independent of nonprobabilistic explanatory considerations: all that is required is that the prediction is unlikely to be true unless the hypothesis itself is true.⁵⁷ Highly specific predictions are always less probable than less specific ones since the former entail the latter. “Precision” here is not the name of an explanatory virtue with a subtle connection to likeliness but rather a new term for a well-known aspect of probabilistic confirmation.

What both of these criticisms tend to show is that IBE cannot be defended by explicating likeliness in terms of loveliness. The problems of parsing out the connection to truth in purely aesthetic terms are too great, and the difficulty of coordinating explanatory virtues and nonexplanatory evidence can be resolved only within a framework that allows the relative strength of diverse considerations to be compared. In fact, what seems to be called for here is a transposition of IBE into a higher and more constructive key that makes full use of the conceptual resources of probability.⁵⁸

Design Reasoning as Inverse Inference

The Basic Conception of Bayesian Inference

The central components of Bayesianism are simple to state yet enormously powerful.⁵⁹ Begin with an assignment of numerical probabilities,

⁵⁷ Barnes, “Inference to the Loveliest Explanation,” 260–1.

⁵⁸ With apologies to van Fraassen for stealing this metaphor from *Laws and Symmetry*, 150. The growing interest in rapprochement between IBE and Bayesian confirmation theory is well illustrated in the exchange between Wesley Salmon and Peter Lipton in *Explanation: Theoretical Approaches and Applications*, ed. Hon and Rakover (Dordrecht: Kluwer, 2001). The second edition of Lipton’s *Inference to the Best Explanation* (New York: Routledge, 2004) gives a great deal of space to Bayesian considerations, by contrast with the first edition where they were entirely absent. For further exploration of a Bayesian recovery of explanatory virtues, see Timothy McGrew, “Confirmation, Heuristics, and Explanatory Reasoning,” *British Journal for the Philosophy of Science* 54 (2003): 553–67.

⁵⁹ For the purposes of this brief sketch I will ignore complications and qualifications that will be of interest only to those already aware of them. A more complete (but still accessible) outline can be found in Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach* (La Salle, IL: Open Court, 1989), 13–38.

drawn from the interval $[0, 1]$, to all propositions. The intuitive idea is to assign lower numbers to less plausible propositions (and zero to outright contradictions) while reserving high numbers for very plausible ones. The assignment covers the conjunction and disjunction of propositions, and in terms of these latter the conditional probability of a proposition A, given B, is defined as the probability of $(A \& B)$ divided by the probability of B, expressions commonly written thus:

$$P(A|B) = P(A \& B) / P(B).$$

One constraint on the assignments is that they must obey the standard axioms of probability theory which tell us, for example, that if $P(H) = k$, then $P(\sim H) = (1 - k)$. It may be difficult for any mere mortal to satisfy this condition of probabilistic coherence across the full range of his beliefs, but Bayesians maintain that this is not in principle any more damaging than the fact that it is difficult to satisfy the condition of deductive consistency among one's beliefs. Both probability theory and deductive logic are normative, not descriptive, and our cognitive faults should not be an excuse for throwing out the standard in either case.

So far there is no mention of inference, but with one's probability assignments in place there is an intuitive method of "updating" probabilities—that is, changing some of them in order to reflect the impact of new information. An elementary consequence of the axioms of probability known as Bayes's Theorem states, in simple form, the following equality:

$$P(H|E) = P(H) \times P(E|H) / P(E).$$

Here $P(H)$ is called the prior probability of H (roughly and intuitively, the probability H has on one's total evidence independent of one's knowledge that E), $P(E|H)$ is the likelihood (roughly, the extent to which one expects E on the supposition that H is true), $P(E)$ is the expectedness of E, and $P(H|E)$ is the posterior probability of H given E, a probability calculable given exact values for the other three parameters. The simplest updating rule tells us that upon learning evidence E with certainty, we should take as our new probability for the hypothesis H the conditional probability $P(H|E)$, so that

$$P_{\text{new}}(H) = P_{\text{old}}(H|E).$$

More sophisticated updating rules are available for situations where our probability for E changes without our becoming certain that E. The process of updating is sometimes called "inverse inference," since it moves from observations to hypotheses rather than, as the hypothetico-deductive method enjoins, from hypotheses to observations.

A moment's reflection on Bayes's Theorem reveals some very attractive features that are of special interest to those seeking to put IBE on a more rigorous footing. All else being equal, a low prior probability of H will be reflected in a low posterior probability. But this seems to capture just the intuition we had with respect to conspiracy theories: their likelihoods notwithstanding, they are nonstarters because they are so incredibly implausible. In Bayesian terms, the low prior reflects that initial implausibility.

We can see the same intuition at work in a comment of Gosset's regarding the hypothesis that a hand of cards is not dealt at random.

I can conceive of circumstances, such as for example dealing a hand of 13 trumps after careful shuffling *by myself*, in which almost any degree of improbability would fail to shake my belief in the hypothesis that the sample was in fact a reasonably random one from a given population.⁶⁰

Gosset's point, in Bayesian terms, is simply that after shuffling the deck carefully himself he would give to the hypothesis R that the cards were randomly ordered a probability very nearly equal to one. As a consequence, the prior probability of nonrandomness $P(\sim R)$ would be very nearly zero; and one can see easily from Bayes's Theorem that it would take an enormously high ratio of the likelihood of nonrandomness to the expectedness of the result, that is, of

$$\frac{P(E|\sim R)}{P(E)},$$

to raise $P(\sim R|E)$ to a value significantly greater than zero.

Once one acquires the knack of manipulating the relevant components of Bayes's Theorem, one can easily give a precise formulation of all manner of intuitions regarding scientific and everyday reasoning. Take the notion that a really surprising experimental outcome gives more confirmation to a theory that predicted it than does an outcome we all expected anyway. In Bayesian terms, this drops straight out of the fact that the denominator on the right hand of Bayes's Theorem is lower in the former case than in the latter: all else being equal, lower expectedness yields higher posterior probabilities. Here we have a ready-made rationale for the attractiveness, at least *prima facie*, of a theory that passes the initial abductive screening with respect to a highly surprising observation.

⁶⁰ Gosset, in Pearson, "'Student' as Statistician," 243.

The literature abounds with Bayesian accounts of the value of diverse evidence⁶¹ and of severe tests,⁶² the virtue of simplicity,⁶³ and the decreasing marginal value of further evidence.⁶⁴ Bayesians have even tackled long-standing problems of the philosophy of science, offering a resolution of Hempel's paradox of the ravens,⁶⁵ a diagnosis of *ad hoc* theory modifications and a solution to the Duhem-Quine problem. It would be difficult to find anything proposed as an explanatory virtue by proponents of IBE that has not been given a plausible Bayesian gloss.

Objectivity and the Problem of the Priors: Subjective Bayesianism

The apparent power of Bayesian methods has won them a considerable following, particularly among philosophers of science who find the resolution of longstanding paradoxes of confirmation in Bayesian terms nearly irresistible. The power and simplicity of standard Bayesian methods is, however, purchased at the price of an extremely strong assumption: that all propositions have point-valued numerical probabilities. No doubt many interesting consequences follow once we have such values in place. But where do the numbers come from?

Perhaps the single most common answer given by those who march under the Bayesian banner is that they are simply a function of an individual's subjective preferences. If I choose to assign a probability of .99 to the hypothesis that I will get heads on the next flip of an apparently fair coin, then there is nothing, according to Subjective Bayesianism, to stop me—provided that I maintain coherence by assigning .01 to the denial of this hypothesis and am careful in my personal choices of related probabilities not to violate the other axioms of probability theory. Since the initial

⁶¹ See Paul Horwich, *Probability and Evidence* (Cambridge: Cambridge University Press, 1982), 118–22, and Howson and Urbach, *Scientific Reasoning*, 84.

⁶² Geoffrey Hellman, "Bayes and Beyond," *Philosophy of Science* 64 (1997): 191–221.

⁶³ See Roger Rosenkrantz, *Inference, Method and Decision: Towards a Bayesian Philosophy of Science* (Dordrecht: D. Reidel, 1977), 93–117. For an alternative approach, compatible with but not motivated by Bayesian considerations, see Mary Hesse, *The Structure of Scientific Inference* (Berkeley: University of California Press, 1974), 223–57. The concept of simplicity is notoriously difficult to unpack in a univocal way, and neither Rosenkrantz nor Hesse is attempting to cover all of its possible meanings. Some Subjective Bayesians have doubts about the viability of any objective analysis of the concept. See Howson and Urbach, *Scientific Reasoning*, 290–2.

⁶⁴ Howson and Urbach, *Scientific Reasoning*, 82–3.

⁶⁵ A particularly important early paper on this, followed more or less closely in many subsequent Bayesian discussions of the raven paradox, is Janina Hosiasson-Lindenbaum, "On Confirmation," *Journal of Symbolic Logic* 5 (1940): 133–68.

probabilities I assign need not be arrived at by any process of inference, there is no problem of accessibility.

On the face of it this is a flagrant violation of objectivity. But Subjective Bayesians reply that one's initial probabilities really do not matter very much since in the end conditionalization will produce epistemic consensus as by an invisible hand. In particular, suppose that two individuals begin with personal prior probability assignments of .99 and .02 for the hypothesis that a given coin is fair, the alternative being that it is biased 3 to 1 in favor of heads. If they consistently maintain that the likelihood of heads (and that of tails) on the hypothesis of fairness is .5 and the likelihood of heads on the bias hypothesis is .75, it can be demonstrated that a sufficiently long series of successive updatings on *any* observed sequence of heads and tails will bring their posterior probabilities for fairness arbitrarily close.

This reply can, however, be inverted. Given any amount of data, it is possible for two subjects to disagree so sharply on the prior probabilities that conditionalizing on the data will still leave their posterior probabilities arbitrarily far apart, so that the convergence of opinions will only occur on the assumption that the priors were not too far apart to begin with. The Subjectivists might respond that such extreme probability assignments would almost never be selected by any two reasonable people. But we have every right to retort that Subjective Bayesians have, by design, provided no resources for judging one personal probability more reasonable than another.⁶⁶

The Subjectivist's answer to the problem of the priors runs headlong into a collision with the objectivity criterion for rational reconstruction. If there is a rational way to frame design reasoning within the bounds of Bayesianism, it will have to be Bayesianism of a nonsubjective type.

Accessibility and Realistic Bayesianism

As so often happens in methodological investigations, our attempts to avoid one difficulty raise another in its place. Subjectivism is unacceptable on grounds of objectivity, but it did supply a ready answer to the question we first raised about the generic Bayesian program: where do the initial probabilities come from? Recall Dembski's complaint (and he is echoing

⁶⁶ The "swamping" argument is pervasive in the Bayesian literature, e.g., Harold Jeffreys, *Scientific Inference* (Cambridge: Cambridge University Press, 1937), 31–4. Horwich highlights the weakness of this response in *Probability and Evidence*, 34–6.

numerous others⁶⁷) that the prior probabilities are often impossible to justify. If this objection can be sustained, then the beautiful probabilistic machinery of Bayesianism cannot be set in motion. The accessibility of the relevant probabilities has become an urgent question.

Without pretending to cover the multitudinous versions of nonsubjective Bayesianism (Objective Bayesianism,⁶⁸ Therapeutic Bayesianism,⁶⁹ Compromised Bayesianism,⁷⁰ Robust Bayesianism,⁷¹ and others too numerous to mention), I want to suggest a modest and realistic way of meeting the challenge of accessibility while retaining some of the benefits of the Bayesian approach. We should begin, I suggest, with a precise model and explore the confirmational relations that hold in that model among idealized data and hypotheses with sharp probabilities to see how the parameters covary. The basic conception of Bayesianism gives us precisely what we are looking for at this stage, provided that we do not fall into the trap of mistaking the model for the state of our knowledge in actual reasoning. The confirmational relations stand out clearly, and in many cases the model really does seem to supply a rationale for our intuitions and a resolution for classic puzzles about scientific inference. Having established this about the model, we can relax the unrealistic requirements of precision and assess the confirmation of real hypotheses in terms of the suboptimal status of our data, representing our knowledge in many cases with a confidence interval rather than with a point value. Frequently, enough structure is retained to enable us to make very definite judgments.⁷²

⁶⁷ Among many others, see Henry E. Kyburg, Jr., *The Logical Foundations of Statistical Inference* (Dordrecht: D. Reidel, 1974) and *Epistemology and Inference* (Minneapolis: University of Minnesota Press, 1983). See also the beginning of Isaac Levi's article "Ignorance, Probability, and Rational Choice," *Synthese* 53 (1982): 387–417, and the remarks from E. T. Jaynes quoted in Roger Rosenkrantz, *Inference, Method and Decision*, 53.

⁶⁸ Best represented by Roger Rosenkrantz, *Inference, Method and Decision*.

⁶⁹ Paul Horwich, *Probability and Evidence*.

⁷⁰ Wesley Salmon, "Rationality and Objectivity in Science, or Tom Kuhn Meets Tom Bayes," in *Scientific Theories*, ed. C. Wade Savage, Minnesota Studies in the Philosophy of Science, vol. 14 (Minneapolis: University of Minnesota Press, 1990).

⁷¹ Peter Walley, "Inferences from Multinomial Data: Learning about a Bag of Marbles," *Journal of the Royal Statistical Society, Series B (Methodological)* 58 (1996): 3–57.

⁷² This idea is not idiosyncratic. For similar approaches see Kyburg, *Epistemology and Inference*, 146–9, and Walley, "Inferences from Multinomial Data." For a survey of the mathematical representation of probabilities when the strong ordering constraints operative in point-valued Bayesian approaches are relaxed, see particularly section 3 of Peter Fishburn, "The Axioms of Subjective Probability," *Statistical Science* 1 (1986): 335–58. Note that by "subjective" Fishburn means roughly "epistemic" or "internal" as contrasted with external definitions of probabilities in terms of frequencies in sets of events.

The process can be illustrated by returning to the Caputo case. Let D be the design hypothesis that Caputo cheated and E be the evidence that the actual pattern was \mathcal{C} . Writing out Bayes's Theorem we have

$$P(D|E) = P(D) \times P(E|D) / P(E).$$

Expanding the denominator according to the theorem on total probability gives us

$$P(D|E) = P(D) \times P(E|D) / [P(D) \times P(E|D)] + [P(\sim D) \times P(E|\sim D)],$$

an equation that looks less complicated when one realizes that the right side is a simple fraction of the form $p/(p + q)$.

Realistically, none of us can justify the selection of one precise number for most of the terms on the right side of the Bayesian equation. But it does not follow that the equation gives us no rational guidance in drawing a conclusion about Caputo. For if, as is surely true, we are in a position to judge reasonably that $P(E|D) > P(E)$, then we will know that E counts in favor of D : we will know which direction the evidence points. And if we can make a rational estimate of the value of the fraction on the right side of Bayes's Theorem, we can go further and estimate the posterior probability as well.

How can we make this judgment? In a wide range of interesting cases, straight frequencies provide the means for evaluating some of the parameters and statistical calculation yields the rest. We know from experience that bizarrely skewed outcomes are common when people cheat (cheating at cards is proverbial), and this provides us with information to use in estimating $P(E|D)$. And the background information provided by Dembski, together with our knowledge of human motives and political allegiances, does not leave us in doubt about the direction in which the outcome would be skewed if Caputo were to cheat.

By contrast, when people use fair randomizing methods (such as an urn), the odds of a skewed outcome on chance are readily calculated and become lower the more skewed the outcome. In the Caputo case we will therefore use precisely the same value for $P(E|\sim D)$ that Dembski uses in his eliminative design inference.⁷³ Even a rough estimate based on our informal sample of the percentage of local political officials who are corrupt enough to cheat for their parties will enable us to generate an approximate value of $P(D)$ which we can represent with an interval of probabilities rather than with a single number. And although the vagueness of our estimates will carry over into a measure of vagueness in the value of the posterior, the very

⁷³ Here for simplicity I follow Dembski in taking urn models to be an unproblematic source of random distributions. See *The Design Inference*, 11.

serious disproportion between the modest size of $P(D) \times P(E|D)$ and the vanishingly small value of $P(\sim D) \times P(E|\sim D)$ assures us that we can be off by an order of magnitude with each of our frequency estimates and still have an enormously high probability that Caputo cheated.

This analysis enables us to see why a growing disproportion in favor of Democrats provides an increasingly strong case against Caputo, a fundamental test for the sensitivity of the reconstruction of this inference. As the preponderance of Democrats increases with each successive assignment, the likelihood factor $P(E|\sim D)$ shrinks, and in an indefinitely long series of Democrats it will shrink toward zero. In consequence, the posterior probability will approach one, since that is the limit of the fraction $p/(p + q)$ where p is a positive constant and q approaches zero.

What impact does the vagueness of my frequency-based estimates have on the cogency of my reasoning? I estimate from my limited experience of the world that as many as 2 percent of all politicians are venal enough to cheat in a way that will produce visible results, but I might be wrong in either direction: the proportion is not likely to be higher than 20 percent or lower than one fifth of a percent given the sample I have, but those are fairly wide margins. Something similar goes for the likelihood $P(E|D)$: it would be excessively skeptical, not to say counterintuitive, to maintain that I cannot get even a rough order-of-magnitude fix on this probability, but an order of magnitude leaves a fair amount of room for error. The upshot is that my conviction that Caputo is cheating as I see his ballot assignments from one season to the next should grow slowly, starting (given my estimate of the frequency of cheating in general) from something negligible, moving through a faint suspicion to become real doubt as the probabilities work their way through the region where the variance of my samples could make the inference swing either way, and finally crystallizing into increasing confidence as the runaway odds against this result on chance outstrip the margins of error of my original estimates.

It might be considered a drawback of this analysis that it can leave us uncertain in the intermediate period where the evidence is radically inconclusive. But I think on the contrary that this demonstrates the power of the Realistic Bayesian approach to model design reasoning. Our inferences cannot be precise when our evidence is vague, and they cannot be decisive when the evidence is seriously mixed. It seems intuitively wrong to insist that there is a single ratio of Democrats to Republicans that should trigger a design inference automatically. And our intuitive reactions as we play with the parameters of the Caputo case suggest that there is a range of outcomes that would leave us indecisive.

What about the supplementation of our background information about Caputo that raised questions about the sensitivity of Dembski's model? As

we suggested when introducing the additional factors, the Bayesian framework allows us to represent the new information very naturally. Knowledge that Caputo is a stern pietist reduces our prior probability for cheating, thereby reducing the posterior probability. Information depressing the likelihood of a high proportion of Democrats given that Caputo tries to cheat will similarly undermine the inference. In this respect, the probabilistic reconstruction accommodates our intuitions as well as IBE and brings the relevance of each factor into even sharper focus.

How Much Is Lost? Objections and Replies

It is fair to ask whether anything is lost by comparison to the standard (unrealistic) version of Bayesianism when we acknowledge realistic constraints. The answer is that some things are, but once again this turns out to be a virtue of the realistic approach. Realistic Bayesianism leaves open the possibility that in certain cases we may not be in a position to evaluate the likelihood of E on each logically possible hypothesis. When we do not have a partition—an exclusive and exhaustive set of hypotheses—then the likelihood factor $P(E|\sim D)$ may be neither calculable nor estimable from frequencies. In extremely unfavorable circumstances of this sort we might not even be able to estimate a posterior probability for D, and we would therefore be unable rationally to decide whether it is reasonable to believe.

A realistic shrug seems called for here. There are numerous cases both in science and in everyday life where our incomplete evidence leaves us uncertain what to believe. The fact that we can model this in terms of ignorance about parameters that are critical for the inference is actually a sign of the explanatory power of this reconstruction: it illuminates our inability to draw a definite conclusion by pinpointing the parameter for which we lack relevant data. And this in turn may provide us with some helpful guidance by indicating the sort of information that would resolve our uncertainty.

Yet a critic may urge that this answer is insufficient. For it is far more often the case than otherwise that we do not have a partition: and yet we seem to have reasonable confidence in our theories despite missing part of the information Bayes's Theorem requires us to have. We know, for example, that the probability that comets approaching our sun will travel in orbits that approximate conic sections, given that Einstein's theory of General Relativity is true (and some other background assumptions), is close to one: but what is the probability that they will do so given that Einstein's theory is false? We have no obvious grounds for estimating that parameter, and at best very uncertain grounds for estimating it to be much lower than the likelihood on the supposition that Einstein is right. Would it follow on Bayesian

grounds that we must suspend belief regarding the truth of General Relativity?

This is a serious consideration, but its force is mitigated in three ways. First, we have already agreed that realism demands a move away from the existence of precise values across all propositions. It is not terribly surprising that this should include some of the propositions whose probability we would like to know, since the availability of relevant data is not a function of our preferences and interests. Second, in the case of design inferences we are sometimes in a strong position to divide the space of hypotheses into two neat parts, since the absence of design leaves us with a chance hypothesis for which likelihoods can be calculated directly (as in the Caputo scenario) or, at worst, estimated empirically.

The third mitigating consideration is that we may be able to extract useful guidance even from data that are insufficient to yield a posterior probability. Suppose, for example, that $P(E|\sim H_1)$ is not accessible because the available well-defined hypotheses H_1, \dots, H_n do not form a partition. We may still be able to assess the relative merits of H_1 and a particular rival H_2 . For by writing out Bayes's Theorem for each hypothesis and taking their ratio, we cancel out the denominator for both expressions and thereby eliminate the troublesome term, leaving only:

$$\frac{P(H_1|E)}{P(H_2|E)} = \frac{P(H_1) \times P(E|H_1)}{P(H_2) \times P(E|H_2)}$$

It is true that the information is now too meager to give us a definite posterior probability for either H_1 or H_2 , but as Shimony points out the right side of this equation resolves neatly into a ratio of their priors and a ratio of their likelihoods, and if we have information relevant to those ratios we can estimate the relative probabilities of the two hypotheses on such data as we have.⁷⁴

As we saw earlier, Dembski suggests in his criticism of Bayesian approaches that the prior probabilities as well are often impossible to justify, and this would undermine our comparison of posterior probabilities. This seems correct as a criticism of unrealistically precise versions of Bayesianism. But it is less telling against the more moderate version in which we can estimate some priors reasonably from experience, as in the Caputo case. Nevertheless, even if we are in particular cases wholly without rational guidance with respect to the priors—or, as is even more frequently the case, if we are unable to reach consensus about the priors—we can still learn something from a consideration of the likelihood ratio

⁷⁴ Abner Shimony, *Search for a Naturalistic World View*, vol. 1, *Scientific Method and Epistemology* (Cambridge: Cambridge University Press, 1993), 275.

$$\frac{P(E|H_1)}{P(E|H_2)}$$

for this tells us which direction the evidence points as between the two hypotheses. This is information on which two parties can agree even if they cannot reach rational agreement about the relative priors of the hypotheses in question. And when the hypotheses partition neatly into design and chance, as they sometimes do, the likelihood ratio

$$\frac{P(E|D)}{P(E|\sim D)}$$

will tell us in an absolute sense whether the evidence confirms or disconfirms the design hypothesis.

This point sheds light on the example of our evidence for General Relativity. Lacking a partition of the space of hypotheses into well-defined alternatives, we are not in a strong position to estimate the factor $P(E|\sim G)$ —the probability of the evidence given that Einstein’s theory is false. But we are in a very good position to compare it to specific rivals; indeed, the test of a theory against a specific rival is a well known feature of scientific inquiry. It is not clear that we are in a strong position to say that General Relativity will *never* be superseded by a better theory, but the vast majority of scientists are quite confident that it is superior to its rivals. And that is exactly what we should expect if we can get a good estimate of particular likelihoods without a strong sense of the relevant priors and without a useful partition. Once again an intended counterexample actually reveals the explanatory power of the Realistic Bayesian position.

Some statisticians have advocated the comparison of likelihoods as the sole method of scientific inference.⁷⁵ As a rule this is a strategic retreat brought on by their conviction that in the cases of interest the Bayesian priors are not available or not sufficiently precise to permit definite evaluation of posterior probabilities. The pure likelihood approach does have the merit of giving us some useful information, since in practice it not infrequently happens that scientists want to know which of two competing theories receives more support from an experimental outcome. In the event that even a modest Bayesianism should prove indefensible, likelihoods provide a tenable fallback. But since I do not share the pure likelihoodists’ sweeping skepticism about the availability of at least interval-valued prior probabilities, I would rather hold the forward position.

⁷⁵ A. W. F. Edwards, *Likelihood*, exp. ed. (Baltimore: Johns Hopkins University Press, 1992); Richard Royall, *Statistical Inference—A Likelihood Paradigm* (London: Chapman and Hall, 1997).

Suspensions may linger that none of these considerations really answers the accessibility challenge. For surely it is conceivable that we should have a strong and widespread intuition that, say, a certain phenomenon exhibits design and yet have no evident way to make even rough rational estimates regarding the relevant components of the inference once we have anatomized it in accordance with Bayes's Theorem. There seems to be a case in point with the "Fine-Tuning Argument." It is difficult (to say the least) to justify any particular probability distributions for the physical constants on either the chance or the design hypotheses, so our hope of obtaining likelihoods for the actual "life-friendly" range seems small.⁷⁶ And the estimates of the prior probabilities among well-informed people vary wildly. Yet a considerable number of people, many of them not theists, find the Fine-Tuning Argument attractive, and some find it utterly compelling. Why should this not be counted as a counterexample to the rational reconstruction offered by Realistic Bayesians?

The most important thing to note about this challenge is that the possibility of a clash between a strong intuition in a specific case and the resources of a proposed reconstruction cannot be ruled out regardless of one's model of design reasoning. Here again we feel the tension between the verisimilitude and the rationality of our reconstructions. The fit between reconstruction and intuition pertains to the former, but the mathematical underpinnings of Bayes's Theorem pertain to the latter: the two concepts remain distinct. And relaxing the demand for point-valued probabilities, though it certainly increases the verisimilitude of Bayesianism, does not by itself close the gap.

In any such clash we have a choice among three different resolutions, and the possibility of each brings with it a methodological precept. First, we may ultimately decide that the reconstruction needs to be adjusted or replaced. Yet if it has proven fruitful, illuminated much that was dark about the structure of other inferences, and has a powerful internal rationale, this step should be reserved until there is something more than a single irritant—a spectrum of anomalous cases that resist repeated attempts at analysis, perhaps, or an alternative reconstruction that can accommodate the successes of the old as well as the new case in question.⁷⁷ Second, we may decide that the intuition must be abandoned. But when that intuition is widespread among informed people, it is never very satisfying to dismiss it *simply* because it

⁷⁶ The case against a particular version of the Fine-Tuning Argument is spelled out in detail in Timothy McGrew, Lydia McGrew, and Eric Vestrup, "Probabilities and the Fine-Tuning Argument: A Sceptical View," *Mind* 110 (2001): 1027–37.

⁷⁷ Baden Powell's "grand maxim" applies here, that "*having once grasped firmly a great principle, we should be satisfied to leave minor difficulties to await their solution.*" See his *Essays on the Spirit of the Inductive Philosophy, the Unity of Worlds, and the Philosophy of Creation* (London, 1855), 94.

conflicts with the reconstruction, since that merely acknowledges the tension without resolving it. It is better if possible to give a diagnosis of why the intuition might seem plausible even if, when it is expanded and reformulated as an inference, it does not turn out to be cogent. Third, it may turn out upon further consideration that the required information is accessible after all—a resolution of the tension that parallels the resolution of certain skeptical arguments. And the existence of this third possibility highlights another sound methodological prescription: we should not be quick to give up either the intuition or the reconstruction if both have substantial support. For in doing so we risk missing that most satisfying form of enlightenment that comes when one discovers that the solution to the problem was available all along.

Recovering the Explanatory Virtues

If we view our incomplete and uncertain evidence from the perspective of a Realistic Bayesian model, we can supply a rational underpinning for those aspects of IBE that seemed at once so attractive and so maddeningly difficult to justify in the face of skeptical challenges. The initial abductive screening, for example, consisted in noting a surprising event and then subsuming it under a hypothesis that made it a matter of course. In probabilistic terms, the surprise is a low expectedness—a low prior probability for E. A hypothesis that removes the surprise is one on which E has a higher probability. But now the function of the abductive screening is clearer: a hypothesis will be abductively attractive just in case and to the extent that the “abductive ratio”

$$\frac{P(E|D)}{P(E|\sim D)}$$

is large, since the increment of change from the prior to the posterior probability of the hypothesis will be sensitive to this ratio. Of course, given the realism of our position, we have to admit that in some cases we might not be able to estimate this ratio with any confidence. Inexpert as I am in matters automotive, I may not be in a position to say whether a bad solenoid would raise the likelihood for my car troubles above their base rate. But it is plausible that those are just the sort of circumstances in which we are not sure whether introducing a hypothesis would remove our surprise.

The probabilistic analysis sharpens our understanding of the explanatory criteria. Peirce sometimes writes as if the hypothesis must actually entail the surprising evidence,⁷⁸ and since E has *ex hypothesi* low expectedness this

⁷⁸ See Peirce, *Collected Papers*, 2.776, where he says rather incautiously that the theory would “render necessary” the conclusion.

strong condition is sufficient to guarantee a high ratio. But it is not necessary. If the expectedness of E is very low, the likelihood factor $P(E|D)$ need be only *relatively* higher in order to produce confirmation, and this is consistent with both being low in absolute terms.

The reformulation of IBE in probabilistic terms also enables us to untangle the skein of problems surrounding the best of a bad lot criticism. A high abductive ratio is no indication of a high prior; and when van Fraassen urges that IBE may, for all we know, leave us sifting through a meager set of hypotheses with little chance of containing the true explanation, the “little chance” he has in mind appears to be a function of the prior probabilities alone. This is a telling criticism of the standard model in which the initial abduction counts as an inference, however weak. For an inference, if it is not fallacious, contributes in and of itself something to our justification for believing the conclusion.

But transposing IBE into a probabilistic key deflects this argument. The abductive ratio does not contain all of the information we need on the right side of Bayes’s Theorem. The relevant question for inference is not whether the prior *per se* or the abductive ratio *per se* is great: each is helpful *ceteris paribus*, but neither is sufficient. We would have no more right to be content with a selection of hypotheses with high priors and abominable abductive ratios than *vice versa*. An initial trawl through the space of potential hypotheses looking for either of these characteristics alone is a *heuristic*, a method of starting the search for an explanation rather than an inferential step in its own right.⁷⁹ And while a particular heuristic may be criticized if it is demonstrably less efficient than an alternative heuristic, it makes no sense to complain that it does not supply what no heuristic ever promised to deliver.

The second criticism of IBE—that the theoretical virtues bore no obvious relation to truth—is likewise parried. Insofar as the explanatory virtues can be explicated in probabilistic terms, they have an intrinsic connection to truth in virtue of their contribution to a high posterior probability. Unlike loveliness, probability is directly connected to truth: for a theory to be probable is for it to be probably true.⁸⁰

⁷⁹ It is not clear that Peirce would have been happy with this interpretation, though this may be due to the fact that it is extremely difficult to reconcile all of Peirce’s own statements on the matter. But it finds some support in his suggestion that the point of abduction is to give rise to “explanations of phenomena held as hopeful suggestions” (*Collected Papers*, 5.196). See Harry Frankfurt, “Peirce’s Notion of Abduction,” *Journal of Philosophy* 55 (1958): 593–7.

⁸⁰ There is, of course, work to be done on the concept of epistemic probability. For some indications of what that concept will look like and how it will need to be defended, see McGrew, “Direct Inference and the Problem of Induction,” *The Monist* 84 (2001): 153–78. The view defended there owes much to that developed in J. M. Keynes, *A Treatise on Probability* and particularly Henry Kyburg, *Probability and the Logic of Rational Belief* (Middletown, CT: Wesleyan University Press, 1961). It is diametrically opposed to the concept advanced by Alvin Plantinga in *Warrant and Proper Function* (Oxford: Oxford University Press, 1993), 159–75.

In our survey of the second criticism of IBE we stopped briefly to note that the explanatory virtue of precision can be given a purely probabilistic basis. The value of precision in prediction, as Barnes notes, is that a precise prediction is unlikely to be true unless the hypothesis itself is true; and this relation holds equally when we explain a phenomenon with great precision rather than predicting it. Within our model this is represented by the fact that all else being equal, the lower the likelihood $P(E|\sim D)$ the higher the posterior of D will be. This elementary fact can be unfolded elegantly to provide a general rationale for the desirability of severe tests.⁸¹

What of Lipton's complaint that reducing loveliness to likeliness would trivialize the project of IBE? The foregoing analysis suggests that this worry was misplaced. IBE is only trivialized if we ascribe no virtues to hypotheses unless they emerge from Bayesian conditionalization with high posterior probabilities. But the rationales given above all suggest, to the contrary, that explanatory virtues should be explicated in terms of other components of Bayes's Theorem.

Curiously, this underscores the very conceptual distinction Lipton was trying to make. The explanatory virtues are all *ceteris paribus* conditions, and explanatory loveliness in any particular respect is neither necessary nor sufficient for likeliness. But only a probabilistic analysis of the theoretical virtues enables us to see clearly why they are subject to a *ceteris paribus* qualification, how they relate to one another, and how their joint exemplification can, after all, underwrite the likeliness of our hypotheses.

The probabilistic reformulation of IBE retains its comparative nature, as we have seen above, and it helps to resolve a curious tension between two descriptions of it. Thagard insists that the inference is comparative; but Lipton suggests that the introduction of a lovely rival need not reduce the loveliness of a hypothesis already proposed. But now we can see clearly why this is so. Because the explanatory virtues are reformulated in *ceteris paribus* terms, a rival that is lovely *in one sense*—say, having a high likelihood—will not reduce the comparable loveliness of the hypothesis in question. But the posterior probability, or “likeliness,” of a hypothesis H_i will depend on the sum of the weighted likelihoods of its rivals, where the weighting factors are the priors of those rivals:

$$P(H_1) \times P(E|H_1) + \dots + P(H_n) \times P(E|H_n)$$

for this is just the expanded denominator of Bayes's Theorem when the hypotheses H_1, \dots, H_n form a partition. Any term in this sum that is large by comparison to the weighted likelihood of the hypothesis of interest will reduce the posterior probability of H_i . And this explains why, to put

⁸¹ See Geoffrey Hellman, “Bayes and Beyond,” 197–9.

Thagard's point in Bayesian terms, the inference to a posterior probability is irreducibly comparative.

Conclusion

With the recovery of a significant number of the explanatory virtues, the Realistic Bayesian model falls heir to the desirable features of the nonprobabilistic account of IBE as a framework for the reconstruction of design inferences. The inferential mechanism is topic neutral, and in view of the extraordinary range of arguments that can be represented in this fashion, the reconstruction of design reasoning within a Realistic Bayesian framework provides extensive integration.

Even so, it may not be the last word in reconstructing design reasoning. The ineliminable tension between verisimilitude and rationality means that the project of rational reconstruction is open-ended, since we can never be quite sure that we have achieved a perfect balance between the requirements of reason and our inarticulate intuitions. But in comparison to each of the alternatives considered, the Realistic Bayesian model exhibits tangible superiority. It is sensitive to known relevant factors on which some rival reconstructions founder, and it incorporates the likelihood on chance that plays such a prominent role in Dembski's reconstruction. The shift from point-valued probabilities to more realistic representations of our knowledge is a great improvement over the artificiality of Subjective Bayesianism in terms of both objectivity and accessibility. And reformulating IBE within the probabilistic framework addresses the questions of relevance and sensitivity that nonprobabilistic IBE, for all its charms, lacks the internal resources to answer. For all of these reasons it is a model worthy of serious attention and development by anyone interested in giving a rational reconstruction of design inferences.