17

Inference to the Best or Only Explanation

17.1. CAN INFERENCE TO BEST OR ONLY EXPLANATION BE RATIONAL WITHOUT THE PSR?

A number of scientific inference schemes have the structure: We have a phenomenon Φ – ideally, one that comprises a cluster of different phenomena – and a number of explanatory hypotheses $H_1, \ldots, H_n$, that is, hypotheses each of which would explain Φ were the hypothesis in question true. Of these, H_i is the best or most plausible or most probable. Hence, probably H_i holds. The question of how we tell which of the hypotheses should be chosen – what "best" or "most plausible" means – is of course the area where much controversy lies, but I will avoid this question. Schemes having this general form I will label *inference to best explanation* (IBE), though some Bayesian schemes not always contained under that label will qualify, as long as the hypotheses among which one is choosing on Bayesian grounds are *explanatory*.

I will be interested in this inference scheme to the extent that it is considered to be truth-directed in a realist sense. The nonrealist about inference to best explanation will thus not be interested in what follows. Indeed, some may find the following considerations moving them to this nonrealist stance.

IBE then holds that given a number of putative explanations, the one that is in some sense best is likely true. But what if in fact the phenomenon has no explanation at all? What if the phenomenon is just a brute fact? This possibility would not, of course, itself be an explanatory hypothesis. Thus, it would not fall within the purview of IBE at all. I want to suggest that it is a possibility for which the IBE defender needs to account. Even if one solved the difficult problem of deciding which of a set of competing

explanations is best and why bestness is a sign of truth, the problem of why we should think there is an explanation at all would remain. The threat of the no-explanation hypothesis is *in addition to* the well-known threat that an unknown hypothesis might explain the phenomenon. Unlike an unknown hypothesis, the no-explanation hypothesis is a hypothesis that is worked out in as much detail as any of the particular explanatory hypotheses, but it is simpler. It just states that the phenomenon happened for no reason at all.

Obviously, this is not a worry we have in everyday doxastic practices. If we find that the airplane that crashed had a metal-fatigue failure in the ailerons that would produce a crash in, say, 5 percent of flights, and if we find no other putative explanation, then we simply conclude that probably the metal-fatigue failure caused the crash. The idea that given such a simple and clear theory about how the crash could have happened nonetheless it might be that there was no cause strikes us as even more bizarre than that there should be no cause in a case in which there is no theory about how the crash could have happened.

But if the PSR fails to be necessarily true, are we right to dismiss this possibility? Why should we not consider the hypothesis of *no explanation* in addition to the explanatory hypotheses? It seems to be simpler than the other hypotheses. The probability of this hypothesis is, moreover, inscrutable, and indeed perhaps objectively the hypothesis has no probability – neither high nor low, neither 1 nor 0 – since it posits an anomic event and one might think that objective probabilities can only be assigned in the presence of laws. In other words, the possibility that the PSR might be false seems to introduce a defeater to IBE-based inferences.

There may seem to be something circular about first having argued for the metaphysical necessity of the PSR in Chapter 16 as the best explanation of why the PSR holds in everyday situations, and then arguing in this section that IBE itself depends on the PSR. This accusation would be fair if we were trying to justify IBE. But that is not the point. It is a given that IBE is a good inferential truth-directed practice. Using this practice, we can argue for the metaphysical necessity of the PSR. Furthermore, however, we can consider the argument that IBE *itself* in part depends for its success on the PSR.

17.2. PREFERENCE FOR EXPLANATORY THEORIES

We tend to prefer explanatory theories to nonexplanatory ones. Take any scientific theory, say, neo-Darwinian evolution. We can try to consider

a weaker theory that is the same theory minus the distinctive causal or explanatory claims that the theory makes. Thus, instead of saying that natural selection and random mutation explain the mammalian eye, the weaker theory will say that there was such and such a sequence of organisms, that after a cosmic ray hit a DNA molecule, it changed in such and such a way, and so forth. The weaker theory is just as useful for all practical purposes, because it produces exactly the same predictions. Moreover, we can retain any counterfactual claims, if we think we need them practically. Thus, we can say that were the glass to be dropped, it would fall, as long as we do not imply that it would fall *because* of the dropping. (An occasionalist will use counterfactual discourse in this way.) Since the weaker theory does the job and fits the data equally well, why not go for it instead?

In general, given two theories that, as far as we know, fit the observed data equally well and give the same amount of predictions, but where one theory claims to explain the data and the other neither claims to explain nor not to explain – take, for instance, the claims of Hellenistic "empiricist" physicians – IBE would tell us to opt for the explanatory theory. This is an instance of use of IBE that cannot simply be reduced to Bayesian epistemology, at least not in the absence of the PSR. Both theories make the observed data equally likely. If Bayesian epistemology is to explain the difference, the difference will have to lie in the priors. But there is no reason to suppose that the explanatory theory should have a higher prior probability than the nonexplanatory theory. In fact, in general, since the nonexplanatory theories make weaker claims than the explanatory ones, they should have higher prior probabilities.

Of course it may be that we only assign a high prior probability to a nonexplanatory theory T when there is an explanatory theory T^* that is strictly stronger than T and whose prior probability is very close to that of T. If so, then confirmation of T would be confirmation of T^*. But *why* should we assign high prior probabilities only to those nonexplanatory theories that have such an explanatory theory associated with them? What justification is there for this? If there is none, then we cannot say that there is anything irrational about those who act differently. (Cf. Section 17.4.2 for probabilistic attempts to answer this question.)

On the other hand, given the PSR, we can give an account of our preference for explanatory theories. There has to be an explanation. Theory A purports to give one and is the best candidate. We say, "Eureka! We have found what we were looking for," and we accept A. On the other

hand, if we had a nonexplanatory theory, we would know that even if that theory is true, we would not be done – there would be a further theory to be given.

Here is a different approach. Suppose, now, that we have a contest between two incompatible theories, *A* and *B*, of comparable complexity, where *A* is explanatory and *B* is not. If we accept *B*, then given the PSR, because *B* is not explanatory, we know that the full story either does not include *B* or goes beyond *B*. Thus, if *B* is correct and a part of the full story (rather than a falsehood or a merely accidental truth), the full story of the phenomenon to be explained is more complex than *B* and hence than *A*. If, on the other hand, *A* is correct, it might well be the full story. Simplicity bids us to prefer *A*, all other considerations being equal. But it does so only given some version of the PSR.

17.3. THE SHERLOCK HOLMES PRINCIPLE

Observe, finally, how strong our commitment is to the PSR, in connection with IBE in the special case in which only one explanation is available. Recall Sherlock Holmes's famous precept "When you have eliminated the impossible, whatever remains, however improbable, must be the truth," which seems at first sight to be a claim about modal logic. But actually the application of this principle uses a version of the PSR, as we see from the continuation of the quote:

> [Holmes:] "We know that he did not come through the door, the window, or the chimney. We also know that he could not have been concealed in the room, as there is no concealment possible. When, then, did he come?"
>
> "He came through the hole in the roof!" I cried.
>
> "Of course he did. He must have done so." (Doyle, 1927, p. 111)

Holmes thus takes the elimination of all other explanations to be tantamount to an elimination of all other possibilities. There is no possibility of the person's coming-to-be ex nihilo in the room.

In our ordinary reasoning, if we are certain that we have enumerated all the possible explanations – as Holmes admittedly had not – and that all but one did not occur, then we are certain that the remaining explanation is the correct one. We do not allow for any possibility of there being no explanation. To modify somewhat an example by David White (1979), suppose that a number of zoo animals were brutally killed

and dismembered. We have absolutely conclusive evidence that no animal had any access to the cages in question, and only one person did. Moreover, we know that no disease or other natural event could have caused the carnage. On the basis of our absolutely conclusive evidence, only that one person could have done it, though there is no direct evidence against him. On the contrary, there is extremely strong character and circumstantial evidence in favor of this man, who was widely reputed to be a holy man who loved animals as St. Francis did. It is highly implausible that he did it. But if it is even more unlikely that anyone or anything else did it, then we will reluctantly conclude that our St. Francis did it.

The only question is whether the evidence ruling out the other possible explanations is solid. Observe, too, that, as Holmes did, no doubt the prosecutor in court will slide between talking of having eliminated all other possible *explanations* and talking of having eliminated all other *possibilities*, and no one will consider as a possibility the option that the dismembering deaths of the animals are a brute uncaused fact. Thus, in our ordinary practice, we assign zero probability to the chance that this was a brute fact. Now, there is no reason why it should be any less likely that we have an unexplained brute fact occur in a case in which there is only one available explanation than in a case in which there is more than one available explanation. Thus, in the latter cases, if we are to be consistent, we should suppose there to be no chance of there failing to be an explanation. The *strength* of this commitment to the PSR also strengthens the claim that the PSR is not itself an inductively confirmed generalization.

Observe that Holmes's principle that if there is only one possible explanation, then this explanation is correct is arguably sufficient as a replacement for the PSR in the most controversial application of the PSR, namely, its application on a global scale to the existence of the cosmos or to the Big Conjunctive Contingent Fact, the conjunction of all contingent true propositions with truth-functional redundancies removed if needed. For as we have seen it can be argued that in those cases the only possible explanation is in terms of the free action of a necessary being. Therefore, by Holmes's principle, this *is* the explanation. This argument was made by David White (1979; see also Katz and Kremer, 1997). And, of course, if we have an explanation of the BCCF, then we have an explanation of every contingent fact, and hence the PSR is true. Thus, Holmes's principle implies the PSR.

17.4.1. Restricted PSRs

17.4.1.a Restriction to Scientific Cases. Let us now consider some alternatives to the PSR that seem to do the job for philosophy of science.

(130) If p is a true proposition that can have a scientific explanation (or, alternately, p is the sort of proposition that can have a scientific explanation), then p has an explanation.

This principle would appear to exclude the BCCF or the proposition that there is a contingent being, since it seems there cannot be a scientific explanation of these facts. Thus, it is a principle that would tell the PSR to keep to its proper place and not lead us into theological realms, grand axiarchic principles, or the like.

However, in fact, (130) does not appear to accomplish the kind of limiting of the explanatory aspirations that it promises. Let S be the set of all contingent physical beings. Suppose that, in fact, all the members of S are discernible, that Leibniz's Principle of Identity of Indiscernibles (PII) holds at least contingently, in the following sense: For each member α of S there is a general nonrelational property P_α, that is, a nonrelational property expressible in general terms without rigidly referring to individuals, places, or times, such that α and only α has P_α. This is a plausible assumption. The argument extends to a more general, and even more plausible, case in which each set of mutually indiscernible contingent beings is of finite cardinality, but I will leave that as an exercise to the reader (cf. Pruss, 2004a).

Consider now the conjunction p of all propositions of the form

(131) There exists a unique contingent being x such that P_α.

Somewhat surprisingly, p is a proposition of a sort that could have a scientific explanation. It is a finite or infinite conjunction of propositions of the form (131). But such propositions *can* have scientific explanations. For instance, if one has a conjunction of seventeen propositions of the form (131), there is no conceptual problem with giving a scientific explanation of why this conjunction is true – one just explains how one of the entities came into existence, how another came into existence, and so on, making sure that one does not involve oneself in a circularity and that one also explains any apparent coincidences. In fact, plausibly, it is logically possible for p to have a scientific explanation, for there could

be a larger universe in which p is still true and where p is explained in terms of the antics of contingent physical entities not falling under any of the descriptions P_α for α in S. In that universe, S would only be a subset of all the contingent physical entities, but this does not affect the argument.

Hence, by (130), p must have an explanation. But, in fact, it seems that p does not have a scientific explanation, for an explanation of p would have to make reference to entities outside S, and no scientific explanation does that, since *in fact* the set S is the set of all contingent physical beings. Thus, p has a nonscientific explanation. Thus, while the motivation for (130) was to clip the wings of the PSR and not allow us to get to nonscientific explanations, (130) fails to accomplish this. Of course, one might deny that the version of the PII used previously is even contingently true as well as denying the more general technical assumption about the finitude of cardinality of each set of mutually indiscernible contingent entities. If so, then the preceding argument would be blocked. But it is highly implausible that it is permissible to infer an explanation of p if and only if this technical assumption holds. That would seem quite ad hoc. And given that (130) fails to give one what one wanted when one proposed it, the simpler full PSR is more plausible.

Now, one might try to block the preceding argument by saying that I misunderstood what (130) is getting at. Granted, p could have a scientific explanation. But it could not have one *in the actual world* where p is in fact a report of the existence of all contingent physical beings. Thus, perhaps, the modality we want in (130) is "possibility in the actual world." But there is no such modality. Unless this modality is going to collapse into ordinary metaphysical possibility, "the actual world" must be a rigid designator. But if w rigidly designates a world and q rigidly designates a proposition, then the proposition *that q holds at w* is either necessary or impossible. Otherwise, we end up with nonsensical claims such as "q holds at w at w_1 but it is not the case that q holds at w at w_2." Hence, if "the actual world" is a rigid designator, whatever can hold at the actual world necessarily holds at the actual world. And so (130) says that p has an explanation if p has a scientific explanation, and that is a mere tautology.

But perhaps the modality in (130) is epistemic. Given that we know p to be the conjunction of *all* propositions of the form (131) for α in S, it is not epistemically possible that p have a scientific explanation. However, epistemic possibility differs from person to person, and so this would mean that p has an explanation provided it is epistemically possible

for at least one person – or maybe one rational person? – that it has a scientific explanation. This is an unacceptably anthropocentric principle. Besides, since views of what can and cannot have a scientific explanation differ widely among rational persons, it would follow in practice that just about any true proposition would have an explanation, except maybe propositions such as *p is true but has no scientific explanation*.

Probably, however, what someone who would propose that the modality in (130) is epistemic really means is that

(132) If *p* is known to be true and it is epistemically possible that *p* has a scientific explanation, then it is rational to believe that *p* has an explanation.

Such a principle would seem to do the job in practice. It allows that if one did not know that, for instance, *p* is the report of the existence of *all* contingent physical beings, then it would be rational to believe that *p* has an explanation, but it might cease to be rational if one learned that *p* actually reports the existence of all contingent physical beings. Likewise, if we learned that some quantum mechanical facts do not have a scientific explanation, then the principle would no longer justify us in thinking that they have an explanation.

There are other variants of the condition in the principle. For instance, one might say (cf. Callender, 2004) that it is rational to believe that *p* has an explanation provided that it is reasonable to think that an explanation would be fruitful, either pragmatically (e.g., by yielding predictions) or explanatorily (in terms of supplying a theory that explains other things). All of these have in common the claim that under some conditions as to what an explanation of *p* would be expected to be like, it is rational to believe *p* has an explanation.

Now, observe that such a principle does not actually reflect the totality of instances in which we are willing to accept the existence of an explanation. Suppose that scientific study demonstrated that some basic constant in the laws of nature when expressed in binary in a physically natural system of units, upon being interpreted as a sequence of 8-bit ASCII codes, spells out a dense five-hundred-page treatise in Danish that has the same combination of being enigmatic and yet insightful as the *Concluding Unscientific Postscript* of Kierkegaard. We would not exactly know what to make of it, just as we do not really know what to make of the *Postscript*. A scientific explanation would not be possible since the constant was assumed to be basic. Since the work is so enigmatic, we could not assume that the claims made in it are straightforwardly true or even accepted by

the author. Besides, the content of the work might be such as not to generate testable predictions even were the content of the work assumed to be all true. Yet, surely, the rational thing to do is to believe that there was an intelligent being that was at least partially responsible for making the constant have the value it has. This does not generate predictions. It might make somewhat plausible the claim that other constants might hide other messages, but even if this claim were not borne out, we would still be reasonable in thinking that there is an explanation in *this* case.

One might think that in this case, one does not need any specific principle for inferring explanations, but simple induction will suffice. Things that look like complex literary works are all written by intelligent agents, we see. This looks like a complex literary work. Hence, it is written by an intelligent agent. But the claim that things that look like complex literary works are written by intelligent agents is no more strongly confirmed empirically than the claim that every contingent proposition has an explanation. Note that in both cases we have cases in which we do not have direct data: there are, after all, complex literary works whose authors are not known and for which we have no independent evidence that they have an author. Furthermore, even if we found out somehow – for instance, because an indisputably divine voice told us – that the mysterious treatise was not composed by an intelligent agent, we would surely at least suspect there is still an explanation and might toss around strange possibilities such as that an intelligent agent was responsible for it in a different way than by composing it (e.g., by designing a nonintelligent machine that produces Kierkegaardian prose), or that somehow the collective unconscious of humankind produced it supernaturally, and so on.

Thus, (132) does not appear sufficient. The restriction to cases in which the epistemically possible explanation is scientific appears ad hoc. It is quite unclear why the epistemic possibility of a scientific explanation gives sufficient evidence for the claim that the proposition has an explanation in a way in which the epistemic possibility of explanation *simpliciter* does not. And of course once one removes the requirement that the explanation be scientific, then the epistemic possibility of the theistic explanation of the BCCF will be enough to secure the reasonableness of thinking that the BCCF has an explanation, and hence that every contingent true proposition has one.

17.4.1.b Detailed Explanation. On the other hand, one might modify (132) by removing the restriction to scientific explanations and inserting

instead a requirement that we have a fairly specific notion of a possible explanation.

(133) If *p* is known to be true and it is epistemically possible that *p* has an explanation that we can spell out in relevant detail, then it is rational to believe that *p* has an explanation.

To block the application of this principle to the BCCF, and hence block the movement from (133) to the rationality of believing the PSR to be at least contingently true, we can insist that the theist does not actually give a putative explanation, but gives a sketch of an explanation. For the theist says that God brought about the BCCF for some reason *R*, but the theist cannot say very much at all about what *R* is. And *something* needs to be said about *R* for the explanation to count as spelled out "in relevant detail" (cf. Grünbaum, 2004).

However, note that this attempt to block the application of (133) fails as a theist *can* give some epistemically possible reasons *R*. She may say, with Rescher's interpretation of Leibniz, that God wanted to optimize a balance of diversity over simplicity. Or she may say, with John Hick, that God wanted to create a world where creation could participate in self-creation on the level of species through evolution and on the individual level through the moral self-improvement and growth of limited persons. Or she might say that God valued the objective good of creatures' praising him freely. There are many possibilities. The theist cannot say *which* of these possible reasons is the one on which God *actually* acted. But (133) does not require that we be antecedently able to say which one of the epistemically possible explanations is the correct one. As long as we can give some sort of a relevantly detailed sketch of *an* epistemically possible explanation, that is all we need.

If one complains that all of these sketches are too vague, consider the similar vagueness that can obtain in putative explanation of why a work of art came into existence. I once saw a number of rocks piled one on top of the other, on a rocky beach in Vancouver. The rocks clearly were purposefully arranged and had some sort of artistic vision. But I could not understand their arrangement. The sketch I could give of a putative agent-based explanation would be quite sketchy indeed: someone arranged them for some artistic purpose or other. But that should be enough for purposes of (133) since this seems a fine application of (133). In fact, to this day, my knowledge of the reasons for this arrangement remains more general than any one of the mentioned proposals in the theistic case. And, in fact, I can offer a theistic proposal that would have a very similar level of specificity:

a highly intelligent agent brought about the BCCF *for some artistic purpose or other.*

Now it is true, I suppose, that if I stared at the work of art, I might eventually come up with possible interpretations that are more specific. But it is implausible that I would need to do that to be justified in simply thinking that there is an explanation. Moreover, some (many?) people do claim to find more and more meaning in the universe as they continue to live.

17.4.2. Probabilistic PSR

17.4.2.a Bayesian Approaches. Perhaps we could simply look at matters in a Bayesian way and assign a very low prior probability to causeless events' occurring. This, too, would give us what we need in daily life and science. But it is not clear what would underwrite such a probability assignment. It is rational to set one's epistemic probabilities in accordance with those objectively assigned by the laws of nature when the laws of nature assign probabilities. But this does not apply in this case, as we have seen in Section 16.4.4. It is perhaps likewise rational to choose one's probabilities in accordance with some principle of indifference. But a principle of indifference is surely not going to tell us that it is unlikely that an object will come into existence *ex nihilo*. It is not clear that any other proposed methods for choosing prior probabilities will help, unless in an ad hoc way they simply *say* that one should assign probabilities in a way that is hostile to things' coming into existence *ex nihilo*.

One might, of course, retreat into the purely subjectivist Bayesian camp. This, however, does not appear satisfactory. It means, for instance, that we cannot say anything about the irrationality of the person who is constantly surprised not to see objects popping into existence *ex nihilo* if her consistent probability assignment included a high probability of their doing so. We *can* criticize such a person on the grounds that her probability assignment was in fact bad for her – that she could not navigate well in the world. But that is not a rationality criticism, just as it is not a rationality criticism to say that it was bad for Fred not to have played the lottery, because the number that he always uses when he plays the lottery had won. It was bad for Fred not to have played the lottery, but there was, we may suppose, nothing irrational about it.

17.4.2.b An Objective Bayesianism. There is, however, an objective Bayesian approach that might help, suggested by Peter Forrest and building on Carnap's original approach. Simply lay down a constraint on prior

probabilities that simpler claims are more probable, by making claims that are expressed in longer sentences in a "natural" language be less probable, where a language is said to be *natural* in this sense if it has simple terms for the more natural universals and in general it is appropriately adapted to the nature of things. Then, unpredictable states of affairs are less likely a priori because predictable states of affairs can be described more briefly: for instance, if a state of affairs is entirely predictable from an earlier state and the laws, to describe the later state of affairs we need only give the laws and the earlier state.

One very natural way to do this is to fix the simple language. Assume the language is written in a finite alphabet of cardinality n. Otherwise one may need nonstandard analysis to handle infinities, but since I am giving this only as an example of the kind of account in question, we do not need to work out the details. Assume that one of the symbols of the alphabet is an "end of sentence" marker. Now, consider a stochastic process that generates strings of symbols, terminating the string as soon as an end of sentence marker is received. This induces a probability measure on a terminated string: the probability of a string is equal to n^{-k} where k is the length including the terminator. Now condition these probabilities on the event E that the generated sentence is grammatical. This yields a probability measure on the set of all sentences that has the property that longer claims are exponentially less likely. This is not quite enough, because if we simply define the probability of a proposition p as the sum of the probabilities of all sentences expressing p, we would probably *not* get a consistent probability measure. For instance, disjunctive claims would likely be less probable than some of their disjuncts. But it is a start. Condition further on the claim that the sentence is a complete descriptor of a world, that is, that one and only one world satisfies the sentence (this can also work in terms of universes, that is, maximal spatiotemporally connected entities, rather than worlds). Now, define the probability of a world as the sum (finite or not) of the probabilities of all the sentences that are complete descriptors of it. And then define the probability of a proposition as the sum (finite or not) of probabilities of worlds at which the proposition holds.

This measure makes simpler worlds more probable than less simple ones. This seems to mean that we have reason to accept simpler laws over more complex ones, and this is exactly what we want for scientific inference. The proposal seems to make many cases come out right. Moreover, it can be argued that any proposal that makes inductive claims come out the way we think they should would assign probabilities in roughly

the way described. If we prefer to model epistemic credences via convex families of probability assignments, we can now take all probability assignments of the preceding sort, ranging over all languages of the right sort, and then make the convex hull of these assignments define the a priori credence distribution.

However, it does not appear that any proposal of this sort will yield a correct epistemology of probabilistic laws, because a formal proposal such as this just does not interact well with probabilistic information. Consider two worlds, each of which contains only one basic particle. According to the laws of both worlds, the particle is of type A until time t_0, at which time it changes into a basic particle either of type B or of type C, and the probability of its changing into a particle of type B is 99 percent. In world w_1, the particle changes into a particle of type B. In world w_2, the particle changes into one of type C. The complexities of the complete descriptions of the two worlds are exactly the same. Thus, any formal proposal of this sort will have to assign exactly the same prior probabilities to these two worlds. But clearly world w_1 is the more probable a priori.

The basic difficulty here is that, plausibly, in some worlds there are stochastic laws, and it seems unlikely that in all worlds the probabilities in these should match the formal ones coming from the objective Bayesianism. We need to retain the probabilities induced by these stochastic laws, since we cannot say that it is objectively just as likely that an improbable event should happen as that a probable one should. However, the probabilities that stochastic laws give rise to tend to behave very differently than the probabilities arising from the objective Bayesianism. A sequence of coin flips each of which comes up heads is overwhelmingly more likely on the objective Bayesianism than a particular less orderly sequence, but stochastic laws are going to assign the same probability no matter whether the sequence is orderly or not. We have two very different patterns of probabilistic assignments, and it does not appear we can make them work coherently together.

And the probabilities induced by stochastic laws are hard to dispute. It is very hard to dispute that if p is the proposition that a process that with probability 99 percent produces a particle of type B has occurred, then the probability that a particle of type B exists given p is 99 percent. If we denied this and allowed the probability to be 50 percent, then we would be very hard-pressed to explain why the fact that a particle of type B is produced in most experiments is evidence for p.

On the other hand, determinism appears to produce problems for the objective Bayesian as well. Suppose that the universe is in fact

bidirectionally deterministic, so that the laws and a complete description of the state of the universe at any one time entail a complete description at all other times. While theoretically determinism could be unidirectional, in practice deterministic theories have been bidirectional. Then, the boundary conditions at the time of the Big Bang (these will be conditions prevailing at the time of the singularity if the Big Bang actually occurred at a point in time and limiting conditions otherwise) determine everything that happens now, and similarly the complete state of the universe now determines those boundary conditions.

Let S_0 be a complete description of the boundary conditions and let S_1 be a complete description of the state of the universe now. There does not appear to be any reason to think that, in fact, there is going to be a way of phrasing S_0 in an appropriate language that is going to produce a shorter sentence than any description of S_1. Conceivably, the initial conditions might have been very nice in a way that allows for an elegant description – maybe the energy levels at the time of the Big Bang were integers whereas now they are irrational numbers, and so on. But there is no reason to think so apart from heady metaphysics, such as theism or optimalism. In any case, suppose that somehow or other we learn (maybe an omniscient being tells us) that the present conditions when backtracked to the time of the Big Bang would yield a state whose descriptions are not shorter than those of the present time. The entropy at that time was, of course, lower, but that refers to macrostates, while the complete description in question is that of the details of a microstate. If so, then our objective Bayesianism is going to ascribe the same prior probability to the hypotheses: "The universe started in S_0 and proceeded via laws L" and "The universe started in S_1 and proceeded via laws L." Since all empirical evidence is balanced between the two hypotheses, the objective Bayesianism will not allow us rationally to say that we know that the universe is more than five minutes old. Now, we do not in fact know the details of S_1. But it would be strange to suppose that were we to learn all the details of the present state, and were we to hear from an omniscient being that the initial state cannot be described any more briefly, then we would conclude that there is a 50 percent chance that the universe started just now.

A crucial assumption in this argument was that the probabilities were assigned to microstates. If instead our objective Bayesianism assigned probabilities to macrolevel descriptions, then the lower entropy at the time of the Big Bang would solve the preceding problem. However, it would do this at a cost. First, the line between the macro and the micro seems

arbitrary and subjective. Second, the objective Bayesianism is designed to ensure that simpler hypotheses are preferable, a right and noble aim. But we have this aim just as much with respect to microstates. We do not just want to be able to confirm macrohypotheses of statistical mechanics: we also want to be able to confirm hypotheses about microstates.

Finally, consider the kind of worldview on which the objective Bayesianism is plausible. Being objective, the Bayesianism has to be responsive to the world in some way. But the probabilities that the Bayesianism is based on are obtained by a probability distribution over sentences, admittedly, sentences of a language other than one of ours, but still sentences. Why should a probability distribution that is natural for *sentences* be appropriate for the world? To put it melodramatically: Why should we think that logos rules the cosmos? Such a probability assignment would be plausible if our ontology were a Tractarian one, or if we literally thought that the universe was ruled by a rational being. But apart from such metaphysics, it appears that a more Humean assignment is more appropriate. The proposed objective Bayesian assignment assigns a disproportionately high probability to arrangements in which there is a pattern, since a pattern makes for briefer descriptions. (A sequence of fifty coin flips in which the results alternate can be described very briefly by saying what the first flip was and what the pattern was.)

Intuitively, prior probabilities should be laid down in accordance with some principle of indifference. But a standard problem with principles of indifference is that the method of partitioning the state space affects the outcomes, and hence the state space must be partitioned in a "natural" way, lest bias should enter in. Choosing the partition in a way that reflects a purely human interest in linguistic description would not be acceptable unless one had some prior reason to think that such a partitioning is indeed natural, as it would be if in some appropriate sense logos ruled the universe. Note, of course, that this criticism does not apply to those objective Bayesians who are willing to assume a priori the existence of God.