

The World, the Deceiver, and *The Face in the Frost*
Lydia McGrew
1/20/2018

Introduction: A tantalizingly incomplete solution to the deceiver challenge

In an appendix to *The Foundations of Knowledge*, Timothy McGrew provides the outline of a solution to the problem of the external world.¹ McGrew uses the device of a Ramsey sentence to argue that the probability of the existence of a deceiver who makes it appear that we live in a real external world must be lower than the probability of a real external world itself, because the ontological commitments of the latter hypothesis will always necessarily be greater than those of the former. In the latter hypothesis, we posit a mental state of the deceiver as a cause of each of the apparently real things that seem to exist in the external world, but the deceiver himself also exists as an entity who is not merely the sum of all of these mental states.

McGrew argues, further, that any time we conditionalize on some particular mental state of our own that we normally take to be caused by real objects in the external world, the gap in probability between realism and the deceiver hypothesis grows larger. He bases this argument on the probabilistic fact that if one theory is strictly simpler than some other theory, the confirmation a given piece of evidence affords to the simpler theory is always greater than the confirmation it affords to the more complex theory--the difference between the old probability of the simpler theory and its new probability is always greater than the comparable difference between the old and new probabilities of the more complex theory.² McGrew's argument thus shows, if we take it to be successful, that the prior probability of a deceiver scenario is lower than the prior probability of realism and also that as we gradually conditionalize on more and more everyday evidence, the gap in probability between the two will continue to grow. This set of conclusions would seem to mean (since we have a great deal of sensory evidence that we normally take to be caused by external-world objects) that what seems intuitively right is actually supported by epistemology and probability theory--namely, that realism is enormously more probable than a deceiver scenario based on all of our evidence, despite the fact that (as McGrew has set up the conditions) the two hypotheses are empirically equivalent in the sense of giving identical probability to all of our sensory and memorial evidence.

If this explanation of what McGrew has shown is correct, his argument is of tremendous importance and deserves a good deal more attention than it has received. In point of fact, everything he has argued probabilistically is completely correct. The only difficulty arises when we come to gauging the significance of the results and in particular of the result concerning the growing gap between realism and an empirically equivalent deceiver scenario.³

McGrew makes no distinction in his 1995 discussion between a growth in the absolute value of the gap between the two hypotheses and a growth in the *ratio* of the probabilities of the two hypotheses. This is significant for the interpretation of his result. Intuitively, what it seems

¹ Timothy McGrew, *The Foundations of Knowledge* (Lanham, MD: Littlefield Adams, 1995), pp. 135ff.

² *Ibid.*, pp. 134-6.

³ I am indebted to Timothy McGrew for the criticism that follows of his own 1995 argument.

should happen is that the ratio of the probabilities of realism and the deceiver hypothesis comes to favor realism heavily, so that, if we take it that hypotheses other than either of these are strongly disconfirmed or even flatly ruled out by the evidence, we are in the end strongly justified in believing realism. Realism, it seems, should have a very high posterior probability conditional on all evidence. This, however, does not follow from what McGrew has argued. For suppose that we grant the point that McGrew has argued concerning the prior probabilities. This point, however, could be satisfied by a very weak inequality. Let us suppose then that the prior inequality is of the ratio 51/49 favoring realism. This could be the case even if both prior probabilities were very low, for example, if we were to model those prior probabilities as .00051/.00049. As we conditionalize on further evidence for which the two theories are empirically equivalent, that ratio *does not change*. This follows from their empirical equivalence and holds despite the growth in the absolute gap between them. Repeated conditionalizing upon evidence for which two theories are empirically equivalent simply causes the ratio of the posteriors to approach the ratio of the priors. If all other theories are effectively ruled out by the evidence, the posterior probabilities for realism and the deceiver hypothesis on this model, when all our evidence is taken into account, will actually *be* .51 and .49. The absolute gap between these is, of course, much larger than the absolute gap between .00051 and .00049, but since the ratio is identical, the final result is not nearly as exciting as we could have wished. Speaking colloquially, realism ends up in this model only slightly more probable than not. The growth of the absolute gap is simply not very epistemically important by itself, aside from further information about the ratio of the priors, and McGrew's Ramsey argument about the priors does not, by itself, justify postulating a prior ratio that *strongly* favors realism as opposed to a ratio that favors realism merely to some unspecified extent. Hence, his argument as a whole does not justify us in concluding that the posterior ratio strongly favors realism, either.

It is a significant thing in itself to argue successfully that realism must have a higher prior and hence a higher posterior probability than an empirically equivalent deceiver scenario; we should not minimize the value of that inequality. But it does not by itself justify the very high posterior probability we intuitively give to realism in our daily activities and that philosophers would like to argue for in answer to skepticism. If a probability only slightly above .5 is considered insufficient for justified belief, this argument would not show that belief in realism is justified. Hence, McGrew's solution to the problem of the external world is incomplete.
 eeen the two.

McGrew's answer to the problem of the external world incomplete is his granting the skeptic for the sake of the argument that the only type of deceiver hypothesis we need to consider is one that is *by definition* empirically equivalent to realism. If H' is defined as being empirically equivalent to H *no matter what* evidence comes to light or no matter how much evidence is involved, then empirical reasoning is brought to a halt by treating paradigmatic *ad hocness* as epistemically legitimate. Such a move would not merely favor the skeptic in the realism/anti-realism debate but would also make it impossible to give strongly preferential rankings to different hypotheses within realism. ~~realism argue that realism~~ is strongly favored over anti-realism if rational empirical reasoning is possible at all--that is, if it is not blocked by treating empirical equivalence *a priori* as an insuperable problem.

1. Enshrining *ad hocness*

Elsewhere I have argued that there should be two concepts of *ad hocness*.⁴ One, the formal concept, involves the disconfirmation of an hypothesis by some evidence E because of the necessity to conjoin the hypothesis with an auxiliary that had (in the old probability distribution) low probability conditional on the hypothesis, while the negation of the hypothesis had no similar problem explaining the evidence. The second concept, which is argumentative and informal, involves irrationally thinking or trying to induce others to think that an hypothesis has not been disconfirmed in this way when in fact it has been. Thus all argumentative *ad hocness* is formal *ad hocness*, though not *vice versa*; a reasoner might admit the correct degree of disconfirmation of his preferred theory by the need to conjoin it with a low-probability auxiliary but might rationally continue to believe the theory if it had enough independent evidence in its favor.

A paradigmatic example of argumentative *ad hocness* is the infamous story of dianetics founder L. Ron Hubbard when he attempted to demonstrate the success of his methods with a protégé. He told the audience that the young woman, having gone through his tutelage to become a "clear," could remember every detail of her past life. The demonstration was a dismal failure; the woman could not answer the audience's questions, including a simple question about the color of Hubbard's tie. Desperate to save face, Hubbard hypothesized that by calling her forward with the phrase, "Come out here *now*," he had fixed her in the present and had thus made it impossible for her to demonstrate her remarkable abilities.⁵

Hubbard's epistemic and rhetorical behavior on this occasion was an attempt to induce irrationality in others by refusing to admit disconfirmation. He was asking the audience to *overlook* the serious disconfirmation of his claims by the failure of the experiment, and he was presenting the auxiliary hypothesis concerning the phrase "come out here *now*" as an excuse for his followers to agree that disconfirmation had not occurred.

Such paradigmatic argumentative *ad hocness* is all the worse if applied repeatedly to multiple items of contrary data. Argumentative *ad hocness* indicates that the person "accommodating" the contrary evidence is refusing to admit that the evidence is, indeed, contrary. In that case, he can never be convinced by evidence to change his mind. From a Bayesian perspective this means that he is not properly conditionalizing on the evidence, since the evidence does in fact significantly disconfirm his preferred theory, while he is retaining the same or virtually the same probability for the theory regardless of evidence. The same principle would apply if the subject were admitting only an extremely minimal and non-threatening degree of disconfirmation when, in fact, the theory is significantly disconfirmed. One can on this analysis admit the importance of some sort of ultimate openness to falsification--that is, disconfirmation so great that the proposition should no longer be believed--without adopting the Popperian baggage that can come with the term "falsification."

The Cartesian skeptic argues that we cannot be justified in believing in external-world realism because of the definitional empirical equivalence between a Deceiver scenario (be it a brain in a vat scenario or a Cartesian scenario without the mechanism of brains and vats) and realism. The idea is that, since the deceiver scenario has been defined to give just as high a

⁴ Lydia McGrew, "On Not Counting the Cost: Ad Hocness and Disconfirmation," *Acta Analytica* 29:491-505

⁵ Martin Gardner, *Fads and Fallacies in the Name of Science* (Mineola, NY: Dover, 1957), pp. 270-271.

probability to *all* of our sensory and memorial evidence as that conferred by realism, we cannot decide between them. Even if the Cartesian skeptic were to admit McGrew's point that deceiverism is to some undefined degree less probable than realism (which he probably would not), McGrew's Ramsey argument permits the skeptic to retreat only to the point of saying that realism is *marginally* more probable than not--perhaps as little as something like .500001, which hardly seems adequate for justified belief in any robust sense.

But it is important to realize the enormous power that the Cartesian skeptic is tacitly giving to this notion of definitional empirical equivalence. Empirical equivalence is being made "trumps" in all possible instances. It seems that, in order to argue that realism can never be justified, the Cartesian skeptic has to adopt some principle like this.

Trumping Empirical Equivalence Thesis--For any empirical hypothesis H, we can define a rival hypothesis H' (or a disjunction of such rivals) that entails the falsehood of H and that is *by definition* empirically equivalent to H for all non-deductive evidence E, whether presently possessed or not. Since we cannot justify strongly disparate prior probabilities for H and H', then, given the possibility that such an H' is true, we cannot be justified in believing H.

The claim that we cannot justify strongly disparate prior probabilities might be defended on the grounds that absolute prior probabilities (prior to all empirical evidence) are indefensible or counterintuitive, an issue I will return to in section 6.

Perhaps the skeptic never thought of himself as making any such sweeping claim for *all* empirical hypotheses, but it is difficult to see what principled grounds there could be for making empirical equivalence "trumps" in the case of realism vs. anti-realism while not making it "trumps" in other disputes. If we cannot be justified in believing realism because of the possibility that all of our experiences are caused by a super-powerful deceiver or deceivers, and because this deceiver hypothesis (or class of hypotheses) has been declared to be by definition empirically equivalent to realism for all evidence, why would the same principle not apply to justification in other contexts? For example, suppose that we consider the conditional probability that real dogs exist *if* realism is true-- $P(D|R \ \& \ E)$, where E is the evidence that most ordinary people have concerning the existence of dogs. If we knew with certainty that realism were true, could we *then* be justified in believing that dogs exist, given our usual empirical evidence? If we accept the skeptic's strictures on belief in realism, then by parity of reasoning it seems that we could not. For it would be possible to construct an in-world skeptical hypothesis according to which dogs do not exist (there are only non-living dog robots) and there is a gigantic conspiracy at a high level to convince people that real dogs (mammalian animals who breathe, breed, etc.) exist. One could then declare this hypothesis D' to be by definition empirically equivalent to D for all evidence E that might otherwise be relevant, and in that case one could never be justified in believing D.

What this shows is that the skeptic's approach to the realism/anti-realism debate enshrines argumentative *ad hocness* and treats it as though it were not irrational. For the Trumping Empirical Equivalence Thesis will ultimately mean that H' cannot be disconfirmed. One way to see this point is to envisage an initial set of evidence as being so entirely inexplicable by anything other than H and H' that it rules out everything else--that is, conditionalizing on that initial evidence makes $P(\sim H \ \& \ \sim H') = 0$. In that case, suppose that *more* evidence beyond that

initial set comes in. By that time, H and H' form a partition, and if the evidence is relevant, it will confirm one or the other and hence disconfirm one or the other. But H' has been declared to be by definition empirically equivalent to H *for all possible non-deductive evidence*, so H could never gain probability at the expense of H', so their probabilities would not change. Further evidence would be irrelevant. But it seems clear that it is possible that there could always be more evidence that would be relevant to, say, deciding between

D: Dogs exist

and

D': Dogs do not exist, but dog robots exist along with a conspiracy to make people believe in dogs.

If D' can be simply declared by fiat to be empirically equivalent to D *for all possible evidence*, then, once everything else is ruled out, it becomes impossible to disconfirm either of them and hence, *inter alia*, impossible to disconfirm D'. But whatever evidence we have so far on a particular empirical issue, it seems always to be possible that we could obtain additional relevant evidence. In the case of the existence of dogs, if we have n apparent dog-like experiences thus far, we could always have n + 1 apparent dog-like experiences. Or we could obtain some new type of evidence, such as the opportunity to do a post-mortem examination of what appeared to be a dog and see that it appears to have mammalian organs and tissues. That evidence, then, would have to be treated as not disconfirming D' by an argumentatively *ad hoc* move--by, e.g., adding the auxiliary of further cleverness on the part of the deceivers and acting as though this rendered the evidence non-disconfirmatory. The same is true if the skeptic admits some slight difference in posterior probability between the skeptical and the realistic scenario but asserts that we can never be justified in considering the realistic scenario to be any more probable than that--i.e., that the ratio between the posterior probabilities cannot undergo significant change. Once the skeptical scenario reaches its "floor" posterior probability, further disconfirmation will be treated as impossible on the grounds of the defined empirical equivalence alone.

Aside from the possibility that all options other than H and H' are strictly ruled out by some evidence, it ought to be possible in principle that some future evidence could both confirm H and disconfirm H' at the same time. While this is not necessarily the case if H and H' do not form a partition, it *may* be the case for some particular item of evidence. But Trumping Empirical Equivalence makes this impossible by declaring that H can never be confirmed at the expense of H', no matter what the evidence.

This brings us to another problem with Trumping Empirical Equivalence: The H' envisaged in the "principle" articulated there is not a well-defined hypothesis at all and hence cannot have a real probability. Because the H' is defined in a purely derivative, piggy-backed fashion *vis a vis* H and *vis a vis* all possible evidence, rather than a definite set of evidence, H' does not have well-specified content. It is an indefinite hypothesis. The content of H' might well need to be augmented and hence changed by the addition of further auxiliaries to account for some later evidence that comes to light. In fact, a modification that would help H' to account for one piece of evidence as well as H does could in principle be harmful to the ability of H' to account for some other piece of evidence. Only hypotheses with definite content can have probabilities, either in a prior or in a posterior probability distribution, so the H' described by the

Trumping Principle should not be taken to be a univocal hypothesis with a genuine probability.

It is important not to become confused here by something akin to a scope shift. It may be correct to say, for some particular *stated* evidence E and an hypothesis H, that it is possible to invent a rival H' that entails the falsehood of H and that gives equal probability to E. But it doesn't follow that, for any hypothesis H, there exists some univocal, meaningful rival H' that is empirically equivalent to H *for all possible E*. Even if an H' is gerrymandered to account for some particular set of evidence just as well as H does, it does not follow that *that* H', without having its content changed by elaboration (and being disconfirmed because of that *ad hoc* change), will account for some other evidence as well as H does, nor that H' can be meaningfully *defined* as *always* doing so while continuing to be the same hypothesis. It is therefore possible that an H' that is *to some degree* empirically similar to H could be disconfirmed by further evidence and that H could in the end be highly justified. Indeed, this possibility must be kept open if the course of inquiry is not to be blocked *a priori*.

The Trumping Empirical Equivalence Thesis, then, ultimately makes on-going, accurate conditionalization on empirical evidence impossible even in cases where it seems intuitively obvious that there could be further relevant empirical evidence that should move our probabilities, disconfirming one hypothesis and confirming another. This insistence that empirical reasoning becomes arbitrarily "stuck" at a certain point merely because of empirical equivalence arises from the move of declaring that a rival hypothesis can simply be *deemed* to be empirically equivalent to some H for all evidence. This insistence is a science-stopper in the strict sense, and in general an empirical-rationality-stopper, not merely in the realist/anti-realist debate, but across the board.

We must, then, reject the Trumping Empirical Equivalence Thesis if we are to maintain the possibility of on-going empirical rationality and not wrongly treat argumentative *ad hocness* as rational. Either Trumping Empirical Equivalence is false, or rational empirical justification by non-deductive reasoning is not possible, since the skeptic of any hypothesis H (other than those known to be true *a priori* or by direct introspection) could always refer to the possible truth of some definitionally empirically equivalent H'.

This point opens up important resources for answering the external-world skeptic. Once Trumping Empirical Equivalence is rejected, there is no principled reason to insist that we must dream up a "packed" hypothesis that explains even all of our *current* evidence and to treat *that* rival as, at most, slightly less probable than the realist scenario. For if we were required to do that at *every particular* moment, the anti-rational effect would be exactly the same as treating H' as *by definition* empirically equivalent to H. Since consistency would require us to apply this in all other scenarios as well, we would be obliged to block justification for any hypothesis on the basis of non-deductive evidence, simply because an empirically equivalent rival could be invented for the existing evidence. At every moment, for every new piece or type of evidence that we encountered or brought to mind, we would have to allow the skeptic to come up with a new version of a rival hypothesis, gerrymandered to account for the new evidence, and summarily declare that to be only (at most) slightly less plausible than realism. This would enshrine *ad hocness* and prevent us from properly conditionalizing on new evidence that confirmed realism at the expense of deceiverism. Instead, as I will suggest in the next two sections, we should evaluate and compare realism to a more generic rival.

2. Ad hoc irrationality and in-world deceivers

To see how the concept of additional evidence and on-going disconfirmation can help in answering deceiver scenarios, consider an in-world deceiver claim--the Case of the Clever Burglars. Suppose that I have a paranoid friend visiting me who fears that the CIA follows him and searches the houses where he stays. One morning he tells me that the CIA has been in the house in the middle of the night before. Upon questioning him, I learn that he is not saying this on the basis of anything he heard or saw in the night. Indeed, he does not expect that he would have seen or heard anything had the CIA broken in. The “deceiver” hypothesis in this case is

B Clever and capable CIA burglars who are very good at hiding their traces have broken into the house in the night and searched it.

This hypothesis is intended to be to some degree empirically similar to the claim that no one has broken into the house, since the burglars are supposed to be good at hiding their traces. But does this mean that it is impossible for it to be disconfirmed by any empirical evidence that is well-explained by a normal night? If so, something is wrong. It ought, epistemically, to be possible to disconfirm the hypothesis and, in the process, to confirm the hypothesis that the night was undisturbed--N (for normal)

N Last night was a normal night in which no one broke into the house or searched it.

B and N, of course, do not form a probabilistic partition. But virtually any evidence that confirms N will *ipso facto* disconfirm not only B but also the presence of ordinary burglars without special powers. And a great deal of evidence that disconfirms B would be of a sort that would confirm N.

In such a situation in real life, I would point out to my friend evidence such as

E₁ The dead bolts on the house doors, which lock only from the inside, are all locked.

My friend, given his paranoia, may then tell me that the CIA must have some ingenious device, the details of which he does not know, for locking dead bolts from the outside while making it look as though they have been locked from the inside. But the very need to postulate this auxiliary hypothesis to deal with a particular piece of contrary evidence, while N has no need for any such special auxiliary, explains why B is disconfirmed by E₁, and my friend should admit it. After all, merely saying that the CIA burglars are clever does not in itself mean that they have such a surprisingly ingenious device for mimicking the locking of dead bolts from the inside. The existence of such an ingenious device actually has quite low probability conditional on B. E₁ is quite well-explained by N. Arguably, it is better explained by N than by ~N, so it confirms N.

I might then bring up

E₂ None of the burglar alarms went off in the night, and the burglar alarm is protected by a complicated password, which I never write down, and it is working

properly as confirmed by a test this morning.

Again, the paranoid friend can hypothesize that the CIA has found *some* unknown way of figuring out my password for my burglar alarm system and thus turning it off and back on without trace, but this, too, will be an auxiliary that was fairly improbable on B itself, while $\sim B$ has no such handicap. Hence, my friend should admit that the postulation of this auxiliary to account for E_2 comes with a cost--B has been disconfirmed once again.

And so it goes for various points. I have (let's say) extremely sharp hearing, sleep lightly, and got up several times, and I heard and saw nothing out of the ordinary in the night. The floor is messy and includes my child's careful arrangement of dinosaur models, in front of most of the windows, where nothing is crushed or appears out of place from its position the previous night. There is new-fallen snow outside, which shows no traces of footprints, etc.

The important point is just this: While B implied a *certain* degree, perhaps even a high degree, of cleverness, technological capability, and knowledge on the part of the CIA agents, it need not and should not be taken to be *definitionally* empirically equivalent to the hypothesis that last night was a perfectly normal night at my house in which no one broke in at all. If my friend were to keep on saying, "But I told you already that these are very clever agents who know how to hide their traces" in response to all evidence, refusing to admit that any such evidence could disconfirm B, and basing this refusal on the fact that he had simply *defined* B in such a way as to be immune to such disconfirmation, he would be irrational. Whatever the prior probability of B, its posterior probability certainly should be significantly lower after conditionalizing on a set of evidence such as this.

It would even be possible to elaborate the Clever Burglar hypothesis so as to account for evidence concerning the background probability B. For example, suppose that we confront the paranoid friend with a CIA agent who tells him how, in his past experience, the organization actually searches houses and who asserts (with details) that some trace would have been left in my house had this occurred the previous night. The paranoid friend can simply say that the CIA agent doesn't know about *this* clandestine program that is tracking *him*. Suppose that the paranoid friend admits that he actually has *no* prior evidence of such CIA activity but points out that this is what would be expected if the agents have *always* in the past been sufficiently clever at hiding their traces as to make themselves empirically invisible. Should we at that point, even if not accepting his theory, retreat to saying that it has an equal or near-equal probability to N because, *via* elaboration, he has made it empirically equivalent for all evidence relevant to the *prior* probability as well as the specific evidence concerning my house?

The in-world case shows how enshrining *ad hocness* means not letting the evidence speak. If empirical equivalence were given as much clout as the external-world skeptic gives it, the same principle would apply to in-world cases, since a sufficiently elaborate in-world skeptical scenario can be gerrymandered to be empirically equivalent to a non-skeptical hypothesis for all data up to the present, including data relevant to the prior probability.

3. Generic Deceiverism and Generic Realism

Consider the following statement of realism, which I will call "generic realism":

GR: There exists a fairly stable extramental physical world to which I have fairly reliable sensory and memorial access.

Now consider a statement of the Deceiver hypothesis, which I will call “generic deceiverism”:

GD: Generic realism is false, and there exists a powerful Deceiver (or group of them) who wishes to produce evidence that will cause me to think that there exists a fairly stable extramental physical world to which I have fairly reliable sensory and memorial access.

I propose that, if we reject the Trumping Empirical Equivalence Thesis, we can treat GD and GR as rival hypotheses and see how they do in competition *vis a vis* the empirical evidence that we possess.

GD, of course, is not by definition empirically equivalent to GR. Nor has it even been gerrymandered to be empirically equivalent for all of the relevant evidence that I can think of right now. (Much of this paper will consist in arguing that much of the evidence I can think of is better explained by GR.) But as discussed above, it was only the idea that Deceiverism is *by definition* empirically equivalent to Realism that generated a motivation for trying to “pack” Deceiverism so that it can account for all the evidence we can think of just as well as Realism does. If we reject Trumping Empirical Equivalence, we are free to consider a more generic rival deceiver hypothesis, to call gradually to mind and attend carefully to various items and classes of our evidence, and to see in a natural epistemic fashion the effect of this cumulative case.

The resemblance to the clever burglar hypothesis in the last section should be clear. In both cases we are using an hypothesis involving deception that has *some* explanatory force but is also capable of straightforward disconfirmation. B would have been empirically equivalent to N for *some* evidence. The fact that the house has not been visibly ransacked, for example, will not distinguish between B and N, since presumably clever CIA stalkers who want to hide their traces will not go around trashing houses like common burglars looking for money or saleable goods. But B and N do not account equally well for the deadbolted doors.

In point of fact, GD sounds a lot like the deceiver hypothesis we philosophers have thought of ourselves as discussing all along. The deceiver in GD is said to be powerful, hence to have significant ability to bring his wishes about as well as a desire to make me believe in a real physical world, and the broadly stated definition of GR is pasted in as the state of affairs the deceiver wishes to simulate. GD is defined by the *prima facie* content of the general class of anti-realist deceiver scenarios. In the nature of the case, since such hypotheses involve the concept of a deceiver, GD will have a connection to realism. To some degree the likelihood of evidence given GD will resemble that given GR, but not to an indefinite extent.

By “treating GR and GD as rivals” I mean simply this: I propose to model the evaluation of GR and GD by treating them *as if* they have prior probabilities (prior to all empirical evidence) greater than zero and less than one (without specifying what these prior probabilities are) and *as if* those prior probabilities are at least equal to each other, though retaining the possibility that GR has a higher prior probability than GD due to McGrew's original considerations. I will examine the details of my available phenomenal evidence to see which hypothesis is confirmed by it and to estimate whether that confirmation is weak, strong, very strong, etc.⁶

⁶While I will frequently use plural pronouns such as “our” and “we,” obviously all the evidence I am describing is actually my own evidence.

GR and GD do not form a prior partition. Even assuming the existence of one Cartesian subject, there are possibilities incompatible with either of these hypotheses. It is possible prior to all specific memorial and sensory evidence that there is an external world, that there is no deceiver, but that I do not have even approximately reliable sensory or memorial access to that world. In that case the world might have any of an infinite or near-infinite number of different actual structures, but I would not know much of anything about it. It is possible, prior to all experience, that there is no deceiver, that there is a real external world, and that I have relatively reliable access to it, but that it is so wildly unstable that my experience will be like that of an ongoing chaotic hallucination. It is possible, prior to all experience, that there is no deceiver and no world external to my mind and that my mind is just randomly generating images. And so forth.

There is, however, a reason why generic realism and deceiver scenarios (where the latter include brain-in-vat scenarios) have been so commonly treated as the only players in the epistemic game, and that is quite simply because the denial of both has zero or near-zero likelihood *vis a vis* even a *portion* of the evidence that we in fact have. For example, if I do not have even relatively reliable access to some real world and there is no deceiver, and if the real world is quite different from anything that I seem to experience (for example, if the real world is made entirely of pink clouds), then my memories of apparent interaction with an apparently stable and complex world are entirely unexplained. If the world were wildly unstable (in the absence of a deceiver) and if I were to have relatively reliable access to it, my present memory-type experiences would be quite different from what they actually are.

Since we will be looking gradually at a cumulative case concerning realism, I shall follow philosophical tradition by treating the conjunction (\sim GR & \sim GD) as being ruled out *early on* in such a way that it deserves no further consideration. In other words, we may take it that in our model, if we were to conditionalize on even a portion of our normal sensory and apparent-memory evidence, that conjunction disappears from the picture or comes so close to doing so as not to merit further thought, causing GR and GD to form a partition at that point (or as near as makes no difference).

Another relevant point is that, just as in the Case of the Clever Burglars, evidence that confirms GR can simultaneously disconfirm GD *even if* they do not form a partition. Evidence does not *necessarily* do so, even for mutually exclusive hypotheses (if they do not form a partition), but it *may* do so, just as the evidence of the dead bolts on the door both confirms N and disconfirms B even though N and B do not form a partition. Hence, even if we conceive of ourselves as conditionalizing on all evidence simultaneously, GR can be confirmed at the expense of GD or *vice versa*.

There is plausibly some generosity built into any model that treats the prior probability of GD as equal to rather than less than GR. GD already involves one of the features used by McGrew to argue for a strictly lower probability for a deceiver scenario than for realism--namely, the existence of the deceiver over and above the existence of those of his thoughts and actions that produce the evidence in my mind. I am inclined not to push too hard on this point, however, because McGrew's use of the Ramsey sentence depends on empirical equivalence. The argument for a strictly richer ontology under a deceiver scenario involves taking it that there are the same number of evidence-causing entities according to the two hypotheses and that the deceiver himself is an entity over and above all of these. Since my version of the argument turns crucially on *not* treating GD as, by definition, empirically equivalent to GR, this move is not available to me. That is to say, I am (as the argument below will show) allowing for the possibility of deceivers who do not have *all* of the thoughts and do not engage in all of the

actions necessary to give us the various types and quantities of evidence we actually have. It does seem intuitively that treating GR and GD as if they have equal priors is not taking account of the extra ontological complexity of the deceiver himself in GD, whoever he is (or they are), but I am not attempting to model this point, merely leaving it open as a possibility.

I will argue that, initial appearances notwithstanding, GD is a much, much poorer explanation of the evidence that we have than GR. That is, GD gives much lower probability to the evidence we have than GR does. Therefore, even if we treat GD and GR as if they have equivalent, intermediate prior probabilities "to begin with," if we treat the denial of their disjunction as being effectively zeroed out by the evidence, and if we conditionalize on the evidence we have, in all its quantity and variety, GR is strongly justified and has a very high posterior probability.

4. *The Face in the Frost*

We can begin to get a sense of the likelihood problems that bedevil GD from a charming but little-known fantasy novel by John Bellairs called *The Face in the Frost*.⁷ The story, set in an imaginary land, pits the lovable good wizards Roger Bacon and Prospero against the evil sorcerer Melichus. Melichus has obtained a book of black magic that gives him the power to make what one might call pseudo-things (though the book doesn't give them a name). Pseudo-things are remarkably robust apparitions. They are visible, tangible, smellable, and three-dimensional. While they last, they are subject to intersubjective verification. But they lack both the perdurance and the full detail of the real things they pretend to be.

The reader sees the height of Melichus's power to make pseudo-things when Prospero, temporarily separated from his friend Roger, puts up for the night in what he takes to be the real town of Five Dials.⁸ When Prospero asks an old man for directions to an inn, we get the first hint that something is not quite right:

The old man pointed his crooked cane toward a shadowy side street and worked his jaws a couple of times before speaking... "Well, ye'd have yer best luck at the Card Player. Go down that alley and turn right. Ye'll see the sign. Mern crost brig."

Prospero cupped his ear. "What was that last thing you said?"

The old man looked flustered and shook his head, mumbling.

"S no matter. G'by. Dirks in cairn."

The old man's insertion of nonsense phrases is just the beginning. Prospero finds the conversation at the inn "curiously vague and listless...[E]veryone was...saying the same thing in different ways." When he goes up to bed, carrying a candle, something strikes him as odd about the mirror in the hall.

As uneasiness grows on Prospero in his room, he picks up a small box on the table:

It didn't even rattle. The heart-shaped brass lock plate on the front was smooth to his touch. It had no keyhole. He turned the box over, looking for hidden locks and spring releases, but there was nothing...*Why did that mirror bother him?...*[In the hall, he] fished

⁷John Bellairs, *The Face in the Frost* (New York: Ace Publishing, 1978).

⁸Ibid., pp. 88ff.

his metal matchbox out of an inside pocket and struck a light....He lit [a candle] and tiptoed ...to the place where the mirror hung. Prospero stared and felt a chill pass through his body. The mirror showed nothing--not his face, not his candle, not the wall behind him. All he saw was a black glassy surface.

Fighting down rising fear, Prospero went back upstairs and began to knock on doors, at first softly, then sharply. He tried the doors. Locked. Locked. And locked. Like the box, the doors didn't even rattle. On an impulse, he opened his pocket knife and tried to slide the blade into the space between a door and its jamb. The point struck solid wood, for what looked like a crack was merely a black line.⁹

Eventually, in a flourish of drama, the inn and the entire town melts unpleasantly away, leaving Prospero alone in a field.

The Five Dials scene shows what happens when a deceiver lacks the power, the creativity, or the continued motivation to make his deceptions as detailed, realistic, and purduring as mind-independent things. Though Melichus is quite an impressive sorcerer, the illusion can leave gaps that allow the initially deceived subject to suspect the truth. This class of possibilities will come up repeatedly in the analysis that follows.

The Face in the Frost also helps us to see how evidence for or against realism could be better or worse, stronger or weaker, in any direction, and that our approach to the realist/anti-realist debate needs to have the flexibility to handle this epistemic fact. The evidence for realism is (intuitively) stronger if I have clear vision of detailed objects than if I have only fuzzy vision of indeterminate objects. The evidence for the existence of a deceiver would be stronger if the deceiver occasionally revealed himself by letting something lapse. The same is true in the Case of the Clever Burglars. I would be rationally more inclined to believe my friend's scary story if I found a suspicious object dropped on the floor of the house on the morning in question--say, a small weapon of a kind that no one in the house possesses. This would be *some* positive evidence for his theory and hence better evidence than the complete appearance of normalcy. Conversely, the appearance of normalcy is better evidence for normalcy if the house is hard to break into than if it is easy to break into.

Treating GR and GD as our two candidate hypotheses allows us to model this sort of flexibility. A deceiver who, accidentally or on purpose, reveals his existence is compatible with GD rather than being excluded from GD by definition as a failure. So there could be, in theory, evidence for GD that is better than what we have. Conversely, the evidence we have for realism can vary in strength, as discussed in the next section. Such an admission of variation in the ability of the evidence to decide between the two hypotheses would be impossible if we attempted to make realism and deceiverism empirically equivalent for all evidence.

5. What does the evidence say?

The evidence that favors GR over GD is difficult to lay out not because there is any paucity of it but because there is so much of it, and I make no pretense to be displaying all of it out here; time and space would fail to canvass it thoroughly. Here I will lay out several categories of evidence that favor GR over GD while asking the reader to extrapolate in order to get a sense of the quantity of evidence subsumed under each. In many places, to show that I am bearing in mind

⁹ Ibid., pp. 92-93

that we are not allowed, when describing evidence, to assume that extramental objects, persons, and events *really* exist, I will use the prefix “apparent” to refer to them (e.g., apparent people, apparent objects, etc.). Where I do not do so, the prefix should be taken as read.

5.1 Detail, layers, variety

If you mentally survey what you spontaneously think you know about the world, what you think you have seen under a microscope or even a magnifying glass or through a telescope, and all the books and articles you believe that you have read describing what other people have supposedly seen at levels you have never observed yourself, you will find an amazing wealth of detailed information going up and down the scale from microcosm to macrocosm. Even if we leave subatomic particles out of consideration, we can talk about the structure of the atom, the structure of the cell, and the crystalline structure of sand grains. We have--or think we have--an entire periodic table listing various elements and the ways in which they are distinct. As we move up above the medium-sized goods around us, we think of planets, the sun, other stars, and galaxies. All these things are lodged in our present consciousness as supposed entities and, more importantly, entities within entities, microcosmic structures out of which other structures are built.

At a more mundane level, boxes appear to open, as do doors and books, and a near-infinite variety of objects and words are found inside. A car's hood can (it seems) be opened, exposing an engine, which can in turn be taken apart to reveal the many parts that fit together to make the car run, with a logical structure that explains (to the mechanically minded) why and how it works. Buildings appear to have rooms inside, forests appear to be filled with trees we have not seen at one time but do seem to see at another time when we go to check, and the oceans appear to be filled with life forms that we discover only when we go there or look at videos and photos or read reports from others who claim to have been there. Our own (apparent) earth appears to be, from our individual perspective, filled with vast numbers of places we have never seen but that we find to be there when we travel and to meet our expectations based upon maps and descriptions.

It is important to realize that all of this is vast overkill from the perspective of GD. Let us grant that some objects, even many objects, outside ourselves must appear to exist if the deceiver in GD is to be allowed the power to effect his goals to any interesting extent at all. Nonetheless, there is no reason whatsoever for him to go this far. The deceiver need not make apparent things so complicated, so multi-faceted, so many-layered, to serve his purpose. Why should he bother? He could convince us of the existence of an external world without this lavish display of creativity. And even if he wanted to make a set of deceptions as exciting, interlocked, and multi-layered as what our evidence gives us, why assume that he can? Melichus apparently either can't or does not want to. That is why boxes in Five Dials neither rattle nor open, doors don't either, conversations are dull, and faux mirrors do not reflect the light of faux candles. Melichus apparently just can't keep that many balls in the air all at once. It is much easier for him to make a world that is, so to speak, all surface.

Consider the same matter from the perspective of GR. Certainly, GR could be true while the world itself was much simpler than it presently is. In that sense we get no positive prediction of a highly detailed world out of GR, either. But what we are concerned with is comparative likelihood. Extra detail confirms GR over GD just in case the probability of that degree of detail is greater given GR than given GD. And here there seems to be no contest. For if there is a real

world outside of ourselves and independent of what someone is trying to make us think at the moment, we would expect that it has some features with which we are not presently in contact. The individual subject is not the “measure of all things” if GR is true. And if that world is relatively stable and we have relatively stable and reliable access to it, then we should be able to find out about presently unknown features. They are, in a real sense, waiting there to be discovered. No one has to trick us by making them merely appear to be there for our benefit. It makes sense, moreover, that some parts of the world would be smaller than others and would exist inside of those others, and that we should be able to find which are which and how they are organized when we go to look. Moreover, the details that we discover appear to hang together (for example, in the fact that the underlying properties of matter give rise to the readily visible properties) in a way that makes sense if there are real objects outside ourselves that really have this structure. If there are real objects outside ourselves, we will expect that their underlying structure--*whatever* it is--will not be causally independent of their more obvious structure. The inside of the box will have a certain shape because it is related to the *outside* of the box. The normally invisible structure of matter will *explain* its visible properties, and so forth, because these are real, stable things. The various aspects of reality are objectively related to each other. No one decides by an act of will to make us *seem to see* a lower layer, which must be *made to appear* to relate to a higher layer. The different layers and facets of reality we see appear to be related to each other because they are, in an important sense, parts of the same thing. Thus the comparative likelihood of the appearance of detail, and the logical relations among details, favors GR over GD.

Though some of these categorizations are a bit arbitrary, one may as well file under the category of “detail” the apparent fact that man is able to discover laws governing the behavior of matter. Or so my present memories appear to indicate. Before me I have what appears to be a science book, and when I open it I seem to read about Boyle's law: “If the temperature of a gas remains constant, its volume and pressure are inversely related.”

GD states that the powerful deceiver wants us to believe that we live in a relatively stable physical world. But once again, allowing us to think that we have discovered *laws* by which the various apparent entities in this apparent world operate is going much farther than necessary for this purpose. Simpler regularities would do. Attempted scientific experiments, intended to establish the underlying laws that give us some sort of “why” for the behavior of the matter we think surrounds us could be met by a kind of experimental “white noise,” just as Prospero's knife encounters nothing but wood when he attempts to insert it between door and jamb. If we take solipsism fully seriously (which GD requires that we do, except in a Berkeleyan form), the deceiver has had to invent for me, at this present moment, all of the scientific information that I think I know, all of the apparent other minds who have taught me or written books, all of the laws that I right now think I understand, and all of the beliefs I have about the experimental results that have established those laws.

In contrast, if an external world really does exist and if we really do have relatively reliable sensory and memorial access to it, it at least makes sense to ask why the surface regularities we find do hold. Moreover, any true underlying laws that do explain the behavior of matter will be at least in principle available. That is to say, there is no person who could *choose not* to make them up and make us seem to find them. A deceiver could choose to disappoint inquirers, but Nature must give up her secrets.

The third type of detail evidence that tells in favor of GR over GD is what I might call the van Leeuwenhoek evidence. At first this part of the argument might seem to be a restatement of

the initial argument from detail: Van Leeuwenhoek believed (or my apparent memories of what apparent history books say seem to indicate that he believed) that he was not wasting his time by making an instrument that would allow him to see nature at the level of the microcosm, and he apparently found something there. But while the van Leeuwenhoek argument up to that point is just a restatement of the first point--that the deceiver, to generate our evidence, would have to envisage microcosmic worlds and convince us that we have had access to them--it goes beyond that point. For van Leeuwenhoek (seemed to) make his investigation by using one (apparent) physical object--the microscope--to investigate *other* apparent physical objects (microorganisms, cells, etc.). So in order to make us believe that we have access to the micro-level of nature in the way that we appear to have it, the deceiver also has to engage in yet more unnecessary creativity concerning the apparent properties of glass, the way it can be used to magnify things, and the like. And the same goes *a fortiori* for the electron microscope, for which the deceiver has to get yet more creative concerning one set of apparent physical objects and the way in which they seem to enable us to investigate another. This goes beyond simply making us think that we have some sort of access to the micro-world, which could have been accomplished, for example, by telling us simpler stories about specially endowed individuals with unexplained super-sight (like Superman's x-ray vision) that enables them to see germs. Once again, to give us our actual evidence about the apparent physical world, the deceiver must be ultra-creative and ultra-motivated to endow this entirely fake world with interlocking details quite unnecessary to his general project *qua* deceiver.

Again, we would not necessarily *expect* given GR that one type of matter (glass, for example) would allow us to see another type of matter better. But given GR, a search for such instruments and an attempt to build them makes sense. Given GR, for example, we have our own physically mediated sensory access to the world. We even find that it is sometimes better and sometimes worse and that simple actions using our own bodies make it better (for example, squinting or using Galileo's trick of making a lens with one's fingers). Perhaps we can make for ourselves instruments that mimic this access at its best and that extend it still further. We can at least examine real materials around us to see if any of them seem to serve that purpose. The probability is not particularly *high a priori* that we will be able to do so, even given GR, but it is *much higher* than it is on GD. For according to GD, we will be able to make such (apparent) instruments only if the deceiver has fairly arbitrarily chosen to deceive us into believing that we are able to do so.

Another interesting aspect of the apparent detail in our sensory evidence that tells in favor of GR over GD is the fact that we appear to be able to investigate temporary lapses in the reliability of our senses and to assign predictable causes for them--fatigue, drugs or alcohol, a malfunctioning brain--using the resources of the physical world. Thus the GR hypothesis hangs together even "at the edges," where our sensory and memorial access temporarily ceases to be reliable. These are, for the most part, not utterly *unexplained* lapses into sensory disintegration.

GD, to be sure, might very well lead us to fear that the apparent external world would sometimes "come apart at the seams," like the town at Five Dials, simply because the deceiver could not sustain it or decided not to do so any longer. Occasional sensory unreliability may even be said to be predicted by GD. The problem for GD lies in the apparent *explanations* within what appears to be a coherent and plausible real world that encompasses these lapses. For a deceiver sufficiently interested in his world and sufficiently powerful and creative to tell us convincing and apparently well-supported stories (medical and psychological) about occasional lapses in the fabric of our sensory access would seem to be also sufficiently interested, powerful, and creative

not to allow the lapses to happen in the first place! Melichus either voluntarily or involuntarily lets Five Dials melt. What he does *not* do is produce a smooth-talking apparent-doctor (who doesn't lapse into mumbling and gibberish and does not seem to disappear) to explain to Prospero that the entire Five Dials incident had nothing to do with Melichus but was the result of a drug in his food at the previous town, an apparent set of detectives to investigate the poisoning, and an apparent chemist to produce and explain the drug to the victim.

The “overkill” theme as we compare GR and GD continues when we consider the fact that we appear to have more than one *type* of sensory access to the world and that our various types of apparent sensory access dovetail together. Sight and touch produce correlative sensations of apparent three-dimensionality which can be matched predictably after we have experienced them. Hearing gives us the Doppler effect which can be correlated with sight to show objects apparently coming closer, bearing down upon us, and then moving farther away again. We come to be able to predict certain tastes and smells and to correlate them with what we see or touch (a particular smell with what appears to be a skunk, for example, a particular taste with an object that has both an appearance and a texture that we associate with sugar or a lemon). All of this is unnecessary from the perspective of GD, and, as with all these matters of detail, it is by no means built into GD *ab initio* that the deceiver will have either the desire or the capability to give us the appearance of *multiple* senses with these myriad interactions with one another. Sight alone, for example, would be a good trick to pull off and could lead us to believe what the deceiver wants us to believe.

On GR, we would not necessarily expect to have this many senses, but we would certainly expect that whatever senses we do have (since they give us relatively reliable access to a real and stable world) will work together predictably as we use them to perceive real objects. Our experience of one sense, together with GR, enables us to expect congruence with our experiences with another sense, if we find ourselves having other types of sensory experiences (apparently of that same object) at all. The probabilistic impact of this point will be discussed below.

Similarly, there are far more apparent other minds (persons) all communicating coherently and in mutually reinforcing ways about what they say are the same objects than any deceiver needs to conjure up, nor is there reason to believe that any deceiver will be able to sustain the appearance of this many other persons and of their consistent and detailed interactions both with me and with the other (apparent) objects in the world. If GR is true, of course the objects in the physical world, being extramental, can be expected to be intersubjectively perceptible, so it would be expected that multiple subjects, if they do exist, will be able to interact both with them and with each other.

5.2 Predictability and perdurance

There is no question that our own sensory experiences appear to us to be experiences of perduring and predictable objects. I have an experience at the moment of a white shape in front of me, and I seem to remember not simply that this shape is called a “computer desk” but also that I have had many, many other experiences of and types of interaction with this same object. And the same is true for the whole rest of the world of apparent objects that we seem to remember having encountered or presently seem to be encountering: We remember their being apparently solid and perduring, not simply fading away, and we seem to remember huge numbers of occasions on which our expectation of their perdurance has allowed us to predict

their behavior and to use them in known or learned ways. We seem to ourselves to have gradually built up more and more of such experiences over a period of years from what seem to us to be our earliest memories of young childhood onward.

We have, moreover, experiences of what appear to be reports from other people telling us of their interactions with the physical world, telling us of their own experiences of the same things we now encounter at other times, or telling us about when and how some new object or objects came into existence--for example, the planting of a tree or the building of a new neighborhood. The apparent reports to which we have access seem to indicate that people have found an external world all around them for many thousands of years. It is a world that changes, to be sure, but it changes in ways that make sense and that appear to be explicable in terms of antecedent causes--the actions of other agents, animals, or physical regularities. Trees grow; houses are built. Trees are chopped down; houses fall; rocks are gradually weathered away. Bodies decompose, and from the earth to which they return new plants grow which, in turn, are consumed by creatures and contribute to the growth of their bodies. Or so it seems in the world as I now seem to be in contact with it and to remember it.

While there is no problem at all in accounting for these appearances on GR (given that it includes the stability of the external world), GD yields a much greater degree of causal independence among the various apparent instances of our contact with the physical world. The deceiver must be able and willing to sustain the apparent physical world through moment after moment of apparent time; or, to put the matter the same way, he must be able and willing to give me the mountainous quantity of present apparent memories--memories which appear to have been built up during years of my own life--of the ability to predict the continued existence of objects and of encounters with the same objects. He must also have chosen to give me the huge quantity of apparent reports from others of such encounters with a relatively stable world spanning what are said to be thousands of years of time, including descriptions of many scientific experiments that depend crucially upon the expectation that physical objects will not suddenly fade away and will continue to act in a predictable fashion. In inventing all of this, the deceiver must not grow bored, must not change his mind, must desire to give me all of this, for no other reason than the incredible, supererogatory verisimilitude of his deception.

5.3 Verisimilitude

The concept of verisimilitude is relevant to the issue of likelihoods. As Lawrence BonJour points out, it seems clear that what he calls a "quasi-commonsensical hypothesis" (which resembles GR) is a "relatively adequate explanation of the details of our sensory experience."¹⁰ What BonJour calls a "digital" explanation of our experience (a deceiver hypothesis) explains our experience by "some agent or mechanism that produces experiences in perceivers like us in a way that mimics the experience that we would have if the represented world were actual and we were located in it, even though neither of these things is in fact the case." In other words, the very notion of a deceiver *presupposes* that the evidence we have is well-explained by something like GR. Otherwise, why does a deceiver produce these sensations rather than others? This point is particularly clear when it comes to perdurance and predictability.

One could even go farther: Plausibly, if the skeptic were right in his concept of what

¹⁰ Lawrence BonJour, "Foundationalism and the External World," *Noûs*, Vol. 33, S13, p. 244.

rationality requires (i.e., epistemic neutrality or something extremely close to it concerning realism), it would not be meaningful to say that the deceiver is producing experiences "as if" of a mind-independent external world or experiences that "look like" such a world. According to the skeptic's strictures, a truly rational subject would *not* believe that these appearances look like an external world, for a truly rational subject would, after receiving such experiences, be neutral on whether or not there is such a world. He would take those experiences to look like (at most) realism-or-deceiverism. The deceiver can deceive successfully only if we associate these experiences with a stable, mind-independent world as opposed to associating them with a deceiver, but given skepticism, this association is inexplicable in any objective fashion. Should the skeptic say that we just *happen* to be irrational in such a way that we *find ourselves* thinking that these experiences look like a real world? If the deceiver has so much control over our thought processes, why could he not manipulate us to think that vague experiences of amorphous sensations are like the existence of an external world, thus saving himself a lot of trouble? It seems as though, in order to say that our experiences *really* are *like* a mind-independent external world, the concept of verisimilitude must be given some content such as that such experiences make it *reasonable* to think that there is such a world. But that is precisely what the skeptic denies; denying it is of the essence of his position.

This point about the meaninglessness, given skepticism, of any normative idea that our experiences *look like* an external world (which is why the deceiver produces them) brings us back to the piggy-backing issue discussed earlier. I have already argued that a deceiver hypothesis that is *by definition* empirically equivalent to realism is indefinite in content and hence cannot be a real hypothesis with a real probability. The same parasitism upon realism means that the skeptic's strictures against realism undermine the notion of verisimilitude that is indispensable in describing the deceiver's own activities and motives.

In all the areas discussed so far--levels of detail, apparent ability to discover laws, the use of apparent physical objects to investigate others, appearance of multiple, mutually confirming senses and multiple, mutual confirming reports from subjects, perdurance and predictability--it would be difficult to overestimate the sheer quantity of our data in favor of the external world. Indeed, despite the attempt here to itemize it, it would be easy to underestimate it. It is only by even beginning to think about it, item by item, that we begin to realize how much evidence there truly is and of how great a variety. Every single story that every single person has ever told you about any event whatsoever must be an invention of the deceiver if no external world exists. Everything you have ever read in any history book, as well as the book itself, must be his invention. All the apparent external-world items encountered, in all their apparently multifaceted nature, in all your memories of all your daily activities and in all anecdotes and historical incidents were themselves the mere inventions of the deceiver. Day after day, moment after apparent moment, you are interacting with what appear to you to be extramental objects and are finding that they stand up to the test. Apparent year after apparent year your experience is going on without any sudden appearance that the external world or any of the objects in it have unaccountably wavered and melted away.

5.4 Independence

To understand what the evidence says and to see the above argument most clearly, the issue of independence must be brought to the fore. It may or may not be possible to atomize our bare sensory experiences, as they are experienced by (say) an infant or a drugged person--a patch

of red here, a meaningless sound there. Such atomic pure sensations may have little to no evidential value taken individually in favor of GR as opposed to GD. But once we gain even a little bit of inductive experience, and certainly by the time we are (or seem to ourselves to be) children of several years old, any such sensations are simultaneously accompanied by a wealth of seeming memories. These include memorial data that associate sensations with other types of sensation (sight with expected touch, for example), familiarity, apparent perdurance, what we seem to remember having heard other people say, and so forth. As most of us actually experience it, a single sensation is part of a far more complex composite, synchronic experience, the parts of which are related in various ways.

It is these relationships that give the greatest force to the evidence in favor of GR. As I have shown elsewhere, the force of a set of evidence $\{E_1, \dots, E_n\}$ *vis a vis* an hypothesis H can be expressed as the product of the individual Bayes factors for each of the E_i and a correction factor showing their relative probabilistic dependence given each of H and $\sim H$, as follows:

$$\frac{P(E_1|H)}{P(E_1|\sim H)} \times \dots \times \frac{P(E_n|H)}{P(E_n|\sim H)} \times \frac{P(E_1 \& \dots \& E_n|H)/P(E_1|H) \times \dots \times P(E_n|H)}{P(E_1 \& \dots \& E_n|\sim H)/P(E_1|\sim H) \times \dots \times P(E_n|\sim H)}$$

The term on the right is the correction factor, which I have dubbed a measure of "relative consilience."¹¹ This is a "correction factor" in the sense that it shows how much weaker or stronger the case really is than it would be if all of the items of evidence were completely independent of each other given H and $\sim H$, respectively.

The considerations above indicate that the relative consilience factor for GR *vis a vis* GD is strongly top-heavy--that is, that the evidence items are much more probabilistically dependent given GR than given GD and that the case for realism is therefore stronger than it would be if the items of evidence are merely considered individually.

As already noted, GR and GD do not form a partition, so GD is not identical to $\sim GR$. While the relative consilience ratio can be estimated for two hypotheses H_1 and H_2 that do not form a partition, as can individual Bayes factors, there is no guarantee in that case that a top-heavy Bayes factor or RC factor indicates confirmation of the "winning" hypothesis of the two, in the Bayesian sense that it is confirmed by the evidence relative to its own negation. However, I have addressed this above by suggesting that the catchall, which is the negation of both GR and GD, gives a probability of near zero to even a portion of the evidence we have, much less all of it. If this is correct, then, where GR is the salient hypothesis H,

$$\frac{P(E_1 \& \dots \& E_n|H)/P(E_1|H) \times \dots \times P(E_n|H)}{P(E_1 \& \dots \& E_n|\sim H)/P(E_1|\sim H) \times \dots \times P(E_n|\sim H)} \approx \frac{P(E_1 \& \dots \& E_n|GR)/P(E_1|GR) \times \dots \times P(E_n|GR)}{P(E_1 \& \dots \& E_n|GD)/P(E_1|GD) \times \dots \times P(E_n|GD)}$$

To see in more detail how the RC factor produces confirmation of GR *vis a vis* GD, consider some ways in which items of our evidence are more causally independent given GD

¹¹ Lydia McGrew, "Accounting for Dependence: Relative Consilience as a Correction Factor in Cumulative Case Arguments," *Australasian Journal of Philosophy*. 95:560-572. Lydia McGrew, "Evidential Diversity and the Negation of H: A Probabilistic Account of the Value of Varied Evidence," *Ergo* 3:10.

than given GR. Causal independence, of course, is not the same as probabilistic independence, but when we realize that two occurrences are causally independent, this consideration generally (though not invariably) favors treating them as probabilistically independent.

If GR is true, different features of our experiential and memorial evidence are not all independent either within each (loosely conceived) category we have discussed or between categories. For example, various levels of detail in physical reality are not physically independent of each other; rather, more readily observable levels have their features at least in part *because of* the features of the less readily observable levels, which allows us to infer, in relation to our experiences of regularity, the continuation of the less readily observed physical traits that underwrite the observable traits. The various senses fit together in a regular fashion because they are all in contact with the same relatively stable external reality. The language, gestures, and actions of other people interact smoothly both with physical objects and with each other because the physical objects (including people's own bodies) are real and accessible to all. Our multiple senses appear to confirm the various levels of detail in the objects about us because the objects are real and our senses are really encountering them at various levels.

GD, on the other hand, introduces a much higher degree of independence, even (I am strongly inclined to think) an irreducibly higher degree of independence, among all of these features than does GR because of the on-going possibility that the deceiver's desire or ability could stop at some point. The deceiver must be able to create all of these impressions and must desire to do so, and either his ability or his desire might (for all we can tell from GD itself) fail at almost innumerable points along the way before we reach the level of variety, detail, and interconnectedness that we find in the evidence we actually possess.

One can attempt to reduce the apparent arbitrariness in GD by thinking of the deceiver as conceiving of (and presenting to us) imaginary *wholes*, entire objects (tables, chairs, plants, insects) conceived in great detail and at many levels of reality, whose parts are interrelated in the deceiver's mind in exactly the way we imagine the parts and levels of real objects to be, rather than conceiving the objects merely in terms of their sensorily obvious surface properties. But while this reduces the independence among the parts of such a pseudo-object in any given case, it does so only for that particular object, and it does so, moreover, at the cost of imagining a deceiver able and willing to do so even for that particular object, when going that far is by no means necessary for his goals *qua* deceiver. To imagine a deceiver even *capable* of inventing pseudo-nature and other embodied agents as we seem to find them is to imagine a deceiver at least very nearly omnipotent, omniscient, and omncreative (without being omnibenevolent). And once we have said that, we have still left a significant amount of arbitrariness in the deceiver's decision to use those powers to produce so elaborate a hoax

A couple of miniature examples give an idea of how the relative consilience factor works to confirm GR over GD and also how memorial experiences associating larger numbers of sensory experiences confirm GR over GD. Suppose that I have a visual experience such as I usually associate with an apple. Simultaneously I have a memory experience that I generally associate this visual experience with a kind of fruit that can be touched, that produces a tactile sensation like *that*, and that also produces an olfactory sensation like *that*. I seem to myself to reach out and touch the apple and smell it, and indeed I simultaneously experience the apparent memory of past experiences like this, the expected tactile sensation (and the apparent memory of having just predicted such a sensation), and the expected smell sensation. GR states that a mind-independent external world exists and that I have fairly reliable sensory and memorial access to this external world. Given my visual experience, my apparent memory of a coordinated tactile

and smell experience connected with this type of visual experience, and GR, I have some reason to expect that tactile and smell sensation. Suppose, too, that I have a more general apparent memory that visual sensations tend to be associated with tactile sensations and a more general apparent memory that visible objects sometimes have smells as well. (This is aside from my apparent memory of the *specific* tactile and smell sensations that I associate with the apple-type image.) These items of my experience (where apparent memory is included) are fairly strongly probabilistically dependent given GR, which is to say that the ratio

$$\frac{P(E_1 \& E_2 \& E_3 \& E_4 \& E_5 \& E_6 | GR)}{P(E_1 | GR) \times P(E_2 | GR) \times P(E_3 | GR) \times P(E_4 | GR) \times P(E_5 | GR) \times P(E_6 | GR)}$$

is significantly top-heavy where E_1 , E_2 , and E_3 are the sensations, E_4 is the apparent memory of the connection among the three of them in the past, E_5 is the apparent memory of a general connection between sight and touch, and E_6 is the apparent memory of a general connection between sight and smell. RC is a ratio of ratios, and this is its numerator.

What about the denominator of the RC ratio, showing the unification of these items of evidence by GD? Suppose that one had the visual sensation (E_1) and the apparent memories (E_3 , E_4 , and E_5). Would this provide as much reason for expecting the tactile and smell sensations as does GR? Even if the deceiver desires to convince me that I have relatively reliable access to a mind-independent world, and even if he has bothered to give me these apparent memories correlating my senses (both generally and specifically), his ability to deliver on that expectation might fail. Melichus presumably knows that Five Dials, if it were fully real, would not dissolve, but he lets it dissolve nonetheless. Or the deceiver's attention might wander, or he might change his mind and decide to let my experience be less unified in this particular instance, to allow my experience to be less "like" that of a real, mind-independent world than it would be if my expectations of touch and smell were fulfilled. This is a far more live option given GD than a similar failure is given GR.

Moreover, the argument above suggests that E_4 , E_5 , and E_6 each have *individual* Bayes factors that are top-heavy in favor of GR, since a deceiver might well not bother in the first place to give me all three apparent senses of taste, touch, and smell and to give me apparent memories that coordinate them. In contrast, if I have multiple senses at all given GR, they would plausibly coordinate as part of my reliable sensory access to a real external world. They would not generally be in *conflict* unless there were some explanation such as specific hallucinations, drugs, etc.

Another simplified case illustrate a similar point. Suppose that it seems to me that multiple people who do not know one another and are not all following the same source account tell me that they have seen Stonehenge. Or suppose that I have one apparent in-person account of someone's seeing Stonehenge and one apparent report in a book of a different person's seeing Stonehenge. Obviously this case raises the entire point made above that GR allows me to have reason to believe in the existence of other minds by inductively observing their behavior, while GD does not. But to simplify further, consider for the moment just the two apparent reports, R_1 and R_2 , including the knowledge that they attest to particular contents, and the RC ratio. R_1 and R_2 are not independent of each other *modulo* GR. If GR is true, an apparent report from one person of the existence of Stonehenge should lead me to expect more strongly than I did before that I will receive another such apparent report. GR states that there are relatively stable, mind-independent objects in the world and that I have relatively reliable sensory and memorial access

to them. Therefore, given GR, the first report can give me some reason to believe that *one* of the real things in this relatively stable world of intersubjectively accessible objects is, in fact, Stonehenge.¹² If GD is true, however, there is far less reason to expect more than one report of Stonehenge than there is given GR, as the deceiver might not choose to make it seem to me that more than one person has seen Stonehenge and reported it. Since there is no real Stonehenge to be seen, it is entirely up to the deceiver to decide how much apparent evidence to give me. And of course he must be able to produce multiple apparent reports. Given GD, some other person in the world cannot simply come upon Stonehenge and choose to report it, regardless of the deceiver. The deceiver has to produce everything. So R_1 and R_2 are more independent given GD than given GR, and the RC ratio favors GR.

The power of GR to unify data as compared to GD is an important part of the confirmation of GR.

6. Foundationalism and the problem of the priors

In the end, there is no putting off the problem of the priors indefinitely. If one is not a Bayesian personalist, the question of where prior probabilities come from has a fairly straightforward answer for the great majority of inferences: Prior probabilities for most inferences come from all of one's *other* concrete evidence--other, that is, than the evidence on which one is conditionalizing at the moment or deeming oneself to be conditionalizing for purposes of some given argument. From this perspective the distinction between the justification of prior probabilities and the conditionalization step using likelihoods does not mark a gulf in the nature of justification between subjective and objective probabilities. It is a desire for epistemic or analytic clarity, or perhaps a matter of literally having newly received some evidence at a given point in time, that leads one to group empirical evidence into that which acts as background information and justifies the prior probability and various likelihoods, on the one hand, and that on which one conditionalizes to obtain a posterior, on the other.

But no such analysis will do at what we might call the very beginnings of inference, for those propositions (such as GR) so sweeping and so vital that they serve as background to our other inferences about both everyday matters and high-level scientific theories. But can we say anything definite about objective prior probabilities of GR and GD, where "prior" means "prior to *all* sensory and apparent memorial data"? How are such priors to be understood in relation only to, say, conceptual truths and the direct Cartesian knowledge of the existence of one's own mind--i.e., what is available *other than* the vast quantities of sensory and memorial data that we wish to use to find their posterior probabilities?

A special problem arises here for the classical foundationalist, since that position entails that no proposition has an intermediate probability--a probability between zero and one--except in relation to some other evidence. There are, on this view, no such things as propositions that are intrinsically probable without being intrinsically certain. Nor are there intrinsically improbable propositions. All intermediate probabilities exist in virtue of the relation of the

¹²Depending on other information, R_1 and R_2 may be independent of each other *modulo* the proposition "Stonehenge exists" and its negation. Multiple reports may be independent given a specific proposition and its negation while they are dependent given a broader proposition such as GR.

proposition in question to some other known propositions.¹³ Moreover, the only propositions that can be regarded as foundational are such in virtue of a form of special access the subject has to those propositions, which gives him certainty about them. In this way, the classical foundationalist is distinguished from the moderate foundationalist, the direct realist, and the externalist. Obviously, neither GR nor GD satisfies any criterion for being foundational in such a scheme; the entire discussion here would be unnecessary if either of them did.

Rather surprisingly, though McGrew defends this type of foundationalism, his own argument concerning the differential prior probabilities of realism and deceiverism, discussed at the outset, suggests the contrary. For those prior probabilities *must* be construed in a probability distribution when empirical data is not yet known. That is the point of their being prior probabilities in the realism/deceiverism debate. The mere fact that we are considering the prior probabilities of those propositions because we are ultimately interested in explaining the evidence we now have does not mean that the value of the priors (nor even their comparative value) is set by the empirical evidence. To the contrary, these must be "absolute" priors or "ur-priors," yet McGrew argues by the Ramsey sentence that realism has a higher prior probability than does deceiverism.

While it may be small comfort to the classical foundationalist, it is worth stressing that, if we *did* regard the prior probabilities of GR and GD (and of the catchall) as objectively real, these would not be (by my modeling here) high probabilities. My model has treated GR and GD as if they have prior probabilities that are at least equal to one another (though that of GR might be higher than that of GD), not as if either of them is high in some stronger sense. They are, in fact, both sufficiently complex propositions that it seems quite legitimate to treat both of them, prior to *all* concrete evidence, as having low probabilities. Therefore, allowing intermediate probabilities to exist relative only to conceptual truths would not serve the usual purposes of moderate foundationalism--i.e., justifying our ordinary knowledge aside from inference--since both GR and GD would presumably start out with absolute prior probabilities less than .5. This would merely act as a basis for beginning to conditionalize on all of the sensory and memory-type evidence we actually have. A typical Moorean proposition like "I have a hand" would, on such a view, not have some obvious, intermediate, *high* probability without inference. In the absolute prior distribution, it would have a *low* intermediate probability and would acquire a high probability only relative to foundational evidence of the standard type that the classical foundationalist is interested in--qualia, memory-type experiences, and the like, such as I have been discussing throughout this paper.

The discussion so far suggests something like Richard Swinburne's notion of real, absolute prior probabilities for propositions, ranked according to complexity, so that something like "That is a planet" has an absolute prior probability higher than "That is 387 small planets with a common center of gravity."¹⁴ Given the greater independence of the evidence given GD than given GR, as argued above, an elaborated subhypothesis of GD that gives equal probability to our existing evidence to that of GR would be strongly disfavored in a complexity-sensitive ur-prior distribution, because it is significantly more complex, requiring multiple references to the deceiver's continued ability and desire to produce the evidence.

It is understandable that philosophers might balk at accepting the existence of such a vast ur-prior distribution in which every possible empirical proposition has a specific, objective

¹³ Timothy McGrew, *Foundations of Knowledge*, pp. 60-65, 70-72.

¹⁴ Richard Swinburne, *Epistemic Justification* (Oxford: Clarendon Press), pp. 82-99.

probability relative only to conceptual truths and to whatever a Cartesian subject could know without having any concrete empirical experience. For one thing, it seems like the number of propositions in such a distribution might be infinite, which is counterintuitive. The construct of such a prior distribution seems rather artificial. It implies that there is a real answer to the question of what an ideal subject would say about the probability of a proposition such as GR and millions of other propositions--indeed, that he would be able to give them *specific* probabilities--if he had *no evidence* on the topics in question at all other than considerations of relative simplicity.

Another possibility is to regard the model here as involving a normalization of the likelihood function over a threefold partition--GR, GD, and (\sim GR & \sim GD).¹⁵ If we think of it in this way, we can think of the likelihoods as having maximal effect with no requirement to assume real prior probabilities. *In effect*, this acts like a model in which GR, GD, and the catchall are given equal prior probabilities and in which this prior distribution is maximally sensitive to subsequent data. But in fact one does not assume that there are real probabilities but rather simply turns the likelihood function on the three possibilities into a posterior probability. My contention has been that, if one does this, the catchall will (in essence) disappear and GR will end up with a very high probability and GD with a very low probability. This is somewhat different from the prior model in Swinburnian terms, both because the Swinburnian model treats the ur-prior probabilities as real and also because the Swinburnian prior probability of the catchall would presumably be much higher than that of either GR or GD, given that it contains metaphysically simple possibilities such as that no physical objects exist, no deceiver exists, *and* nothing is making it look like any physical objects exist. That, however, becomes a moot point when the priors are swamped.

A third possibility, rather similar to the second, would be to incorporate simplicity and complexity considerations into a non-normalizable prior distribution that is not a probability distribution at all but that ranks a possibly infinite number of propositions according to their relative complexity. This idea, which remains relatively non-specific, is that such a distribution (as in the case of normalizing the likelihood function) would be maximally sensitive to data and would yield a posterior probability when evidence comes in, in accordance with the relative complexity of the possible explanations.¹⁶

The skeptic will be hard-pressed to argue that any of these ideas is unfair to his position *per se*. Can he argue that GD has a real prior probability and that it should be regarded as *much higher* than that of GR, so that it can handle all the disconfirmation I have argued for while still ending up at least equal to GR? It is extremely difficult to see how one could go about making such a case.

The skeptic's better option would be to attempt to argue that we should throw up our hands in despair over the problem of prior probabilities and that his challenge is unanswerable

¹⁵ The concept of normalizing the likelihood function for some set of evidence and a set of hypotheses that form a partition has been suggested by the objective Bayesian Roger Rosenkrantz, *Foundations and Applications of Inductive Probability* (Atascadero, CA: Ridgeview Publishing Company, 1981), 4.4-13. Rosenkrantz implies that the concept originated with Sir Harold Jeffreys, *Theory of Probability* (Oxford: Clarendon Press, 1948) p. 102, but Jeffreys speaks in terms of uninformative equal priors rather than in terms of normalizing the likelihood function.

¹⁶ Something like this has been suggested by Timothy McGrew, personal communication.

for that reason alone. If we cannot even begin to model a Bayesian or quasi-Bayesian inference by making any sensible model of the situation prior to all concrete evidence, then there seems to be no point in talking about GD as "disconfirmed" or about GR as "highly confirmed" by the evidence.

I stated at the outset that my argument would show that realism is highly confirmed if rational empirical inference is possible at all, and that point comes back into play here. For if we despaired in that way, any inference leading to a high probability on the basis of empirical evidence would become impossible. Nor would this be a result of the fact that we could not then be justified in believing realism *per se*. Such an alleged problem would apply even if realism were assumed to be true. There would be no way to say, given realism, that an elaborate conspiracy theory was highly unlikely and a rival, non-conspiracy theory highly probable, since one could always elaborate the conspiracy theory to the point that it would "account for" all evidence we currently have just as well as some rival, non-conspiracy theory, pushing the issue back to prior probabilities. Then one could declare that, since the problem of the priors is insoluble, there is no way to tell which is correct. Such despair, then, is really just another version of Trumping Empirical Equivalence. If we do not allow rationality to be undermined by Trumping Empirical Equivalence, we can let the evidence speak.

Conclusion: A note on direct realism

Despair of defending the propositions contained in GR on a more fundamental basis creates a significant motivation for the push in epistemology to declare such propositions "basic" or to take them to be intrinsically and directly justified in some way that does not involve even tacit rational inference.

But once we see that the tacit principle on which the skeptic's opposition is based, Trumping Empirical Equivalence, would rule out the rationality of inferences even *after* anti-realism is dismissed, the Moorean or pragmatist has just as much need as the foundationalist to point out that Trumping Empirical Equivalence is irrational, not just pragmatically but epistemically. Once that is understood, the possibility opens up for justifying belief in the external world "from scratch."¹⁷

It is certainly true that we find ourselves believing in the external world and trusting our own senses spontaneously and continuously, and if one takes it that spontaneous belief and overwhelming trust are incompatible with rational inference, one will naturally have a problem with the attempt to justify such beliefs inferentially. But what if we reject the opposition between spontaneity, on the one hand, and rational inference, on the other? What if, instead, we accept a robust notion of "inference" that allows rational inference to be incredibly rapid, spontaneous, and inexplicit--even to take place non-linguistically? In that case, we should be open to the following possibility: These beliefs are so spontaneous, so pervasive, seem so obvious, and act as a background to so much else we believe because they are so heavily *overjustified*--hyperrational, one might say. It might seem surprising that overwhelmingly justified propositions should be believed in such a way that they appear (to some) to be held arationally, yet it is not surprising upon reflection. Where a tacit inference is justified by so great a wealth of

¹⁷ Contrast this argument with the move to pragmatism by Eric Olsson, *Against Coherence* (Oxford: Clarendon Press, 2005), pp. 173ff, after arguing that it is not possible to justify realism "from scratch."

data forming so many interlocking and converging lines of evidence, and where a great quantity of that data first comes to the subject at a very early stage of his existence, it is understandable that he should take the inference for granted, find it nearly impossible to doubt, and find it difficult or impossible to articulate his reasons for it. This need not be a sign of the absence of good reasons; it may be a sign of excessive amounts of evidence.

This view of the matter satisfactorily accounts for the fact that we would consider a person insane if he seriously did doubt the existence of the external world. It accounts for the fact that it seems that we cannot get along in the world--as indeed we cannot--without assuming GR to be true and even without assuming other things, such as the existence of other minds, not contained in GR but readily justified on the basis of GR and of our own specific evidence. To all of these points the classical foundationalist can say this: Yes, you are right. Yes, there is a sense in which we are *compelled* to believe in these things. Yes, it would be insane seriously to doubt them. Yes, we do constantly believe them without stopping to worry, wonder, and tease out our reasons for so doing, much less the metalevel explanation of the cogency of those reasons. Let us go on doing so. And let us go on doing so because what compels us in these beliefs is *reason itself*--reason that has been given so much material to work with, material that is interwoven so beautifully and with such complexity, that the reason suffers from an embarrassment of riches and scarcely knows how to explain the matter to anyone who would doubt.

To us as philosophers is given the odd and off-putting task of creating problems where the ordinary man, understandably enough, sees no problem. But if we do so, let us also take on the task of trying to answer those questions thoroughly. Thus we can be both philosophers and real men living in the real world, seeing no conflict between the two.