

The World, the Deceiver, and *The Face in the Frost*

Lydia McGrew, October, 2010

DRAFT

A tantalizingly incomplete solution to the deceiver challenge

In an appendix to *The Foundations of Knowledge* (1995), Timothy McGrew provides the outline of a solution to the problem of the external world. 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 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 real things that we believe 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 (p. 135-6).

McGrew argues, further, that any time we condition 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 (pp. 134-6). 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 condition 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 the deceiver scenario based on all 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 our sensory 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.¹

McGrew makes no distinction in his 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 informal interpretation of his result. Intuitively, what it

¹I am indebted to [suppressed for blind review] for the criticism of McGrew's argument's incompleteness that follows.

seems 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 model those prior probabilities as .00051/.00049. As we condition on further evidence for which the two theories are empirically equivalent, that ratio *does not change*. This follows from their empirical equivalence and is the case despite the growth in the absolute gap between them. Repeated conditioning 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 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 epistemically important by itself, aside from further information about the ratio of the priors, and McGrew's Ramsey argument about the priors does not justify us in postulating a prior ratio that strongly favors realism as opposed to a ratio that favors realism 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 even 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. Hence, McGrew's solution to the problem of the external world is incomplete.

Paradigmatically irrational *ad hocness*

What makes McGrew's answer to the problem of the external world incomplete, I believe, is his granting the skeptic for the sake of the argument that the only deceiver scenario we need to consider is one that is by definition empirically equivalent to realism. While the problem of the external world is often seen as in its very nature "about" empirical equivalence, to the extent that we are concerned not merely with discussing an interesting example of empirical equivalence but also with deciding what it is rational to believe, the treatment of the competing hypotheses as empirically equivalent is, I will argue, a mistake. To see why, we need to make a digression into the subject of *ad hocness*.

My quarry here is not a definition of *ad hocness*. I do not seek to state a set of necessary and sufficient conditions for a main hypothesis, the postulating of an auxiliary hypothesis, or an argumentative move to be *ad hoc*. I am, however, interested in sufficient conditions for the irrational *ad hocness* of an argumentative move, particularly in certain paradigm cases.

Consider the infamous story about how dianetics founder L. Ron Hubbard attempted to

demonstrate the success of his methods with a protege. He told the audience that the young woman could remember every detail of sense perception from 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. (It does not seem that Hubbard said how long this effect would last nor that he proposed a later test without the alleged hindrance.) (Gardner 1957, pp. 270-71)

What makes Hubbard's move here so obviously irrational? Rather than attempting to analyze this case in hypothetico-deductive terms of a main hypothesis with auxiliaries to generate predictions and make absolute falsification possible, I propose instead a simpler analysis: Hubbard's epistemic and rhetorical behavior on this occasion was irrational and attempted to induce irrationality on the part of others because he was refusing to admit disconfirmation. He was asking the audience (or anyone who learned of his *ad hoc* rescue) simply to *overlook* the serious disconfirmation of his claims provided 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 pretend along with him that disconfirmation had not occurred.

Whether or not the adoption of an addition to one's original hypothesis while continuing to believe the core of the original is irrational depends, *inter alia*, upon the prior probability of the original hypothesis, but even if it is not irrational to retain the core of the hypothesis, disconfirmation will have occurred. A paradigmatically irrational move is the pretense that there has been no disconfirmation.

Suppose that I have known my friend John for many decades and have mountains of detailed evidence to the effect that he is a loving and faithful husband and an honorable man. Suppose, then, that a detractor tells me that he has seen John in a restaurant with a woman who is not his wife and has overheard some ambiguous words (which he reports) that could be easily construed to mean that they are having an affair. It may well not be irrational for me to continue to believe that John is a faithful husband and an honorable man, adding the auxiliary hypothesis that this incident occurred but was misinterpreted by the detractor. But for me to remain rational I must admit that the proposition "John is a faithful husband" has to some degree been disconfirmed in the process.

To see this schematically, consider the fact that only a portion (and a fairly small portion) of the probability space previously occupied by "John is a faithful husband" was occupied by the conjunction "John is a faithful husband who went out to dinner innocently with a woman who was not his wife with whom he had an innocent conversation that was easily misunderstood by a bystander to mean that he and the other woman were having an affair." But the probability of "John is a faithful husband" is now co-extensive with the new probability of that conjunction.

Suppose for the sake of concreteness that the prior probability for

J John is a faithful husband

is .99. Let D be

D John went out to dinner innocently with a woman who was not his wife with whom he had an innocent conversation which was easily misunderstood by a bystander to mean

that he and the other woman were having an affair.

Suppose for simplicity that the $P(E|(J \ \& \ \sim D))$ is zero. Suppose that

$$P(E|(J \ \& \ D)) = P(E|\sim J) > 0$$

Suppose (and this is the crucial point) that, in the probability distribution before I receive the evidence, $P(J \ \& \ D)$ is quite low, perhaps as low as .01. Then $(J \ \& \ D)$ and $\sim J$ will have posterior probabilities in proportion to their priors, giving them equal posterior probabilities, which represents notable disconfirmation for J , from .99 to approximately .5. Of course, we can imagine the posterior probabilities going differently, especially if we make the prior probability for J even higher or change things in some other way so that $P(D|J)$ is greater than $P(\sim J)$. Still, given this set-up of the likelihoods, as long as $P(D \ \& \ J) < P(J)$, J will always suffer at least some disconfirmation from the evidence. And the reason is fairly clear, in epistemic terms: We would expect the evidence more strongly, even much more strongly, if J is false than if J is true.

Under these circumstances, I appear irrational to my interlocutor if I airily dismiss his report with a shrug, saying, "Oh, I'm sure that was just some misunderstanding. I've known John for years." This response, like Hubbard's, gives the impression that I am refusing to admit that any disconfirmation of John's faithfulness has occurred.

This analysis helps to explain the increasing appearance of irrationality when multiple pieces of evidence are successively accommodated by the addition of auxiliary hypotheses with no appearance of lessening confidence in the main hypothesis. Such a response to evidence tends to indicate 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.² One can on this analysis admit the importance of some sort of ultimate openness to "falsification"--that is, to disconfirmation so great that the subject should no longer believe the hypothesis--without adopting the Popperian baggage that can come with that term.

Ad hoc irrationality and in-world deceiver scenarios

To make the scenario more like the deceiver case while making it more tractable than a full Cartesian scenario, let us consider the Case of the Clever Burglars. Suppose that I have a friend staying with 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.

²A Bayesian model for this refusal is that of a person who goes back and treats his prior probabilities as higher than they actually were (retrofitting them) whenever this is necessary to accommodate contrary evidence without admitting disconfirmation and having a lower posterior than the one he desires.

Obviously, this hypothesis is intended to be in some sense or to some degree empirically similar to its negation, because the burglars are supposed to be good at hiding their traces. But does this mean that it is impossible for it to be disconfirmed? If so, something is wrong. It ought, epistemically, to be possible to disconfirm the hypothesis, and it seems like it is possible to disconfirm it if we suppose that my house is the sort that it would be difficult to break into and search without leaving traces. 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, no doubt, would 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 amounts to disconfirmation of B, 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. My friend has brought up this particular auxiliary only to deal with E1, and it would be a paradigmatically irrational *ad hoc* move for him to pretend that B is not disconfirmed by E1.

I might then bring up

E2 None of the burglar alarms went off in the night, but the burglar alarm is protected by a complicated password and is working properly as confirmed by a test this morning.

Again, the paranoid friend can hypothesize that the CIA has found a way of figuring out my password for my burglar alarm system and thus turning it off and back without trace, but this, too, will be an auxiliary with which B was not, to begin with, co-extensive. Hence, my friend should admit that the postulation of this auxiliary to account for E2 comes with a cost--B has been disconfirmed once again.

And so it goes for various points--the messy floor, including my child's careful arrangement of dinosaur models, in front of all the windows, where nothing is crushed or appears out of place from its position the previous night, and so forth.

The important point is just this: While B implied a *certain* 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. The clever agents hypothesized by B do not by definition know how to re-lock dead bolts from the outside, do not automatically know the password for my burglar alarm system, and do not know how my child had his many dinosaur models arranged nor how to set them up just right in front of the window after leaving by the window. That is why it would seem relevant to bring up these points in responding to my friend. And if he 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 lower after conditioning on a series of pieces of evidence such as these.

Letting the evidence speak

One can explain the irrationality of the paradigmatic *ad hoc* move in a number of ways. Formally, one can simply say that my friend presumably had some prior probability or other, perhaps a very high one (which might itself be irrational) for B. And in that case, since E1, etc., are disconfirmatory of B, he should lower his probability from whatever it previously was. This seems straightforward. But there is a more general principle at issue, the principle that we must let the evidence speak--a rephrasing of Pierce's maxim, "Do not block the course of inquiry." It would be possible, though it would ultimately seem self-deceiving, for my friend to say that he merely made a mistake about his prior probability and that, in retrofitting it so that the posterior remains high despite the evidence I bring up, he is merely engaging in better and clearer introspection concerning his prior. As such new "revelations" about how tremendously high his prior probability actually was keep coming to light, coincidentally in connection with my attempt to present counterevidence, I will of course conclude in the end that my friend is utterly irrational on this particular subject--completely closed to evidence.

But in fact, we did not imagine that this was his explanation. Instead, we imagined that he explains his intransigence by saying that he is considering B to be by definition empirically equivalent to

B' Last night was a normal night in which no one came into the house.

In the context of discussing the problem of the external world, it is important to realize that this definitional move, too, is a refusal to let the evidence speak. As we discussed above, it is not actually necessary to take B to be definitionally empirically equivalent to B'. In dialectical terms, to do so is to make B immune to disconfirmation (in favor of B') by the very evidence that, it seems, ought to be capable of disconfirming it. And it is because we would not initially take B to be empirically equivalent in all these detailed ways to B' that we would bring up these pieces of detailed evidence in an argument with the paranoid friend.³

A thoroughgoing Bayesian might respond that the definitional move should not matter, for if B is taken to be empirically equivalent to B', everything will simply turn on the prior probabilities of B and B', so it will all come out in the Bayesian wash. But both epistemically and probabilistically, this is a problematic way of looking at things. When my friend asserts B and I contemplate B originally, there is no reason to take B to be tantamount to a lengthy conjunction including further, highly specific and individually implausible auxiliaries about the CIA's

³There is no parallel here in the fact that, prior to investigation, B' contains various possible sub-hypotheses, some of which are ruled out by further evidence. For example, my guest might not previously have noticed the color of my child's dinosaurs, so B' might at first include both "The dinosaurs are all green" and "The dinosaurs are of various colors," one of which will be ruled out in the course of noticing that the dinosaurs are blocking the windows. It does not follow, however, that B' is disconfirmed by the discovery of particular facts of this sort. For B' to be disconfirmed by them, they would have to be more probable on $\sim B'$ than on B', but this is not the case. Similarly, the discovery that, specifically, Einsteinian physics is more nearly correct than Newtonian physics does not disconfirm realism, for that fact is not more probable on the negation of realism.

knowledge of my password, seemingly superhuman abilities to move about the house without disturbing anything and to put things back precisely where they were before, and the like. And it cannot possibly have been tantamount to a conjunction of an unspecified and indefinitely large set of such abilities as would be necessary to accommodate all evidence that might in principle be brought against it, including evidence presently unknown. The definitional move is hence a sign of epistemic laziness, an attempt without actually considering, listing, or taking into account specific details to “pack” B with all that it needs to withstand disconfirmation for all possible future evidence.

Whatever my friend's original evidence for B, it could not in the nature of the case have supported an indefinitely lengthy conjunction of abilities and motivations on the part of the agents to hide their traces--that is, an indefinite amount of empirical equivalence to B', which would otherwise gain probability at the expense of B. Even if he had evidence for some particular set of abilities on the part of such agents, this evidence could not be extrapolated to cover *all possible* evidence that would make B' appear to be true and that would be difficult for them to fake. Hence, his evidence could not give a known prior probability for a form of B construed as definitionally empirically equivalent to B'.

The situation should therefore be modeled by construing B in the more general terms in which it is actually stated and then seeing how well it does or does not account for the evidence we find in my house the next morning. In this way we let the evidence speak.

Generic Deceiverism and Generic Realism

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

GR: There exists a relatively stable extramental physical world to which I have relatively 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 who wishes to make me think that there exists a relatively stable extramental physical world to which I have relatively reliable sensory and memorial access.

I propose that GD and GR should be treated as rival hypotheses and that we should see what happens when we consider the actual phenomenal evidence that we possess. Which hypothesis, GR or GD, is better confirmed by our internally available evidence which does not presuppose either one to be true?

It is plausible that I will eventually be accused by those sympathetic to skepticism of stacking the deck by treating GR and GD as rivals. Later I will give a number of reasons for the fairness of using these two hypotheses in this way. For now, I will give only one: There is nothing *prima facie* unfair about pitting GD against GR, because GD sounds like the deceiver hypothesis we philosophers have thought of ourselves as discussing all along. In particular, it sounds like this because a) the deceiver in GD is said to be powerful, hence to have at least some ability to bring his wishes about as well as a desire to make me believe in a real physical world, and b) the contents of GR are simply pasted in as the state of affairs the deceiver wishes to

simulate.

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 equal prior probabilities greater than zero and less than one (without specifying what these prior probabilities are) and by looking at the details of my available phenomenal evidence to see which hypothesis is confirmed by the evidence and to estimate whether that confirmation is weak, strong, very strong, etc.⁴ Another way to think of this modeling is as if we are normalizing the likelihoods of the hypotheses vis a vis the evidence.⁵

Someone may point out here that GR and GD do not form a prior partition. That is quite true. It is logically possible prior to all empirical evidence that a black hole universe exists without any finite subjects at all, including me, which would make both GR and GD false. It is logically possible prior to evidence that there is a real 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 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 wildly unstable.

There is, however, a reason why realism and deceiver scenarios have been so commonly treated as the only players in the epistemic game, and that is quite simply because no other hypothesis has anything other than zero or near-zero likelihood for the evidence that we in fact have. If I do not exist, I am not having any of the experiences that I am, in fact, having and to which I have access--the experiences that constitute the evidence to be taken into account. If I have no even relatively reliable access to some real world (perhaps a real world made up entirely of pink clouds) and there is no deceiver, then my memories of apparent interaction with an apparently stable and complex world are, in a sense, unexplained--there is no reason to expect them at all. 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.

I shall therefore follow philosophical tradition in treating the conjunction (\sim GR & \sim GD) as being ruled out by the evidence in such a way that it deserves no further consideration. In other words, we may take it that in our model, when we condition on the evidence, that conjunction disappears from the picture or comes so close to doing so as not to merit further thought.

One more point preliminary point is worth making: There is plausibly some generosity already involved in a 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

⁴While I will frequently use plural pronouns such as “our” and “we,” obviously all the evidence I am describing is actually my own evidence.

⁵The concept of normalizing the likelihood function has been suggested by the objective Bayesian Roger Rosenkrantz (1981, 4.4-13, “Historical notes and additional references for chapter 4”). Rosenkrantz implies that the concept originated with Sir Harold Jeffreys (1948, p. 102), but Jeffreys speaks in terms of uninformative equal priors rather than in terms of normalizing the likelihood function. Both Rosenkrantz and Jeffreys emphasize the importance of a formal way of reflecting prior ignorance and permitting evidence to have its maximal effect.

and above the existence of those of his thoughts and actions that produce the evidence in my mind. I am inclined not to push on this point, however, because McGrew's use of the Ramsey sentence depends crucially 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 in the two hypotheses and that the deceiver himself is then an entity over and above all of these. Since my version of the argument turns crucially on *not* treating my GD as 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 evidence we actually have. It does seem intuitively that treating GR and GD as having equal priors is not taking account of the extra ontological complexity of the deceiver himself in GD, but I am deliberately not attempting to model this point because to do so I would have to weigh it against other considerations--considerations central to my argument as a whole--that give GD a higher prior probability than a deceiver scenario that is tailored to account for all evidence. For my purposes it is sufficiently enlightening to treat the prior probabilities as equal, even if this is generous, and I view my argument as supplementing McGrew's in showing that a truly empirically equivalent deceiver scenario is not only less probable to some unspecified extent than GR but is *much* less probable than GR.

In brief, my modeling of the problem of the external world suggests this: Initial appearances notwithstanding, GD has much, much lower likelihood for the evidence that we have than does GR. Therefore, if we treat GD and GR as if they have equivalent, intermediate prior probabilities, if we treat the denial of their disjunction as having zero likelihood, and if we condition on the evidence that we have, GR is strongly justified and has a very high posterior probability. By the same token, the posterior probability of GD will in the end be equivalent to the posterior probability of a version of GD, if such a version exists, that is empirically equivalent to GR (or as nearly equivalent to GR as it is possible for a deceiver hypothesis to be) for the evidence adduced, but since GR will have "eaten up" all the probability space that GD has lost in the process, the posterior probability of this version of GD will be very low.

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*. 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 we might call pseudo-things (though the book has no such name for them). 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 (Bellairs 1978, pp. 88ff). 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 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 (Bellairs 1978, pp. 92-93).

Eventually, in a flourish of drama, the entire town, including the hostess of the inn, melts unpleasantly away, leaving Prospero standing 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 real things. The illusion leaves gaps that allow the initially deceived subject to suspect the truth. And as there is no guarantee and indeed no particular reason to believe that a deceiver will have all of these characteristics to the extent required to account for the evidence, there is no reason to treat GD as equal in likelihood to GR.

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 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 and it is relevant, the prefix should be taken as read.

Detail

If you 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 yourself observed, 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, 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 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. 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 which we discover only when we go 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 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. That much we can expect given GD. Nonetheless, there is no reason whatsoever for him to go this far, nor is there any reason to take GD to entail that he is capable of going 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 do not 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 simpler for him to make the world more like the world according to logical positivism--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. 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, we would expect that it has features with which we are not presently in contact, for 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 and whatever we are able to discover of it--will not be 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 things and what we are discovering are aspects of reality that are actually related to each other in a mind-independent fashion. No one decides arbitrarily 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 that we find strongly favors GR over GD.

Though some of these categorizations are a bit arbitrary, we 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 seventh grade science book, and when I open it I read about, for example, 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. Simple 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 simply be met by a kind of experimental "white noise," just as Prospero's knife encountered nothing but wood when he attempted to insert it between door and jamb. If we take solipsism fully seriously (which GD requires that we do, except in its 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 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 makes sense to ask why the surface regularities we find do hold. Moreover, any true underlying laws which do explain the behavior of matter will be at least in principle available to us--that is to say, there is no person who could choose not to make them up as merely apparent-laws or not to reveal them to us. 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 the outset 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 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 tell us that we have had access to them--it goes

beyond that point. For van Leeuwenhoek made his investigation, and anyone who has ever (seemed to) use a microscope has (seemed to) make his investigation by using one (apparent) physical object--the microscope--to investigate others (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 *mutatis mutandis* 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 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 (with no apparent mechanism) about specially endowed individuals with unexplained super-sight that enables them to see germs. Once again, to give us our actual, detailed 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 our purposes. 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, according to which we will be able to make such (apparent) instruments only if the deceiver has quite 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 causes for them--fatigue, drugs or alcohol, hallucination--using the resources of the physical world. Thus the GR hypothesis hangs together even "at the edges," where our access temporarily ceases to be reliable. These are, for the most part, not merely unexplained lapses into hallucination and 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 and overshadows 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 and fluent 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. 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, 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 dovetail in many ways as we use them to perceive the same real objects. Our experience of one sense, together with GR, enables us to predict our experience with another sense.

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 produce to be convincing, 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 as a subject and with the other (apparent) objects in the world. If GR is true, of course the objects in the physical world are intersubjectively perceptible, so it would be expected that multiple subjects, if they do exist, will be able to interact with them and with each other via their bodies.

Independence

The issue of independence is obviously playing an important role in these evaluations of relative likelihoods. If GR is true, these features are not all independent either within each (loosely conceived) category we have discussed or between categories. 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. The various senses confirm each other because they are all in contact with the same external reality. The language, gestures, and actions of other people interact smoothly both with physical objects and with each other because the physical objects 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 among all these features. 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) quite readily 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 an 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 the whole of nature as we seem to find it, with all its levels of detail and its laws, plus our own apparent sensory access and our own apparent ability to manipulate it, is to imagine a deceiver at least very nearly omnipotent, omniscient, and omncreative, which was no part of GD *ab initio*. 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. The resemblance here should be quite clear to the CIA agents in the original example, who seem to grow almost without limit in powers and prescience as we adduce more and more evidence against a break-in.

Probabilistically, a set of data provides stronger evidence for a theory H_1 than for a theory H_2 if the pieces of evidence in the set are more strongly positively relevant to one another on H_1 than on H_2 (see McGrew 2003). It is easy to assume that pieces of data should be regarded as independent and, indeed, that an assumption of independence will give us the strongest case for some given theory H .⁶ But this is by no means always the case. Sometimes, given McGrew's unification result, dependence among data favors H where that unification is greater *modulo* H than *modulo* a rival.

We can see how this could work in the case of the external world when we consider how the existence of the external world generally unifies our apparent experiences or reports of a given object within the world. For example, suppose that multiple people who do not know one another tell me that they have seen Stonehenge. The reports do not seem to be independent of each other *modulo* GR. If GR is true, a 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 report, since the first report together with GR gives me reason to believe that one of the real objects in this relatively stable, extramental (hence intersubjectively accessible) world is, in fact, Stonehenge.⁷ If GD is true, however, there is no particular reason to expect more than one report of Stonehenge, as the deceiver might not choose to make it seem to me that more than one person has seen Stonehenge or (if we assume that more than one person actually exists) might not

⁶For example, independent witnesses to an event are often considered to give us the strongest evidence for the occurrence of the event, though there are circumstances under which even this assumption regarding witness testimony may be incorrect (McGrew & McGrew 2009, pp. 631-5).

⁷It is plausible that these testimonies are independent of each other *modulo* the highly specific proposition "Stonehenge exists." It is interesting to consider that the effect of multiple reports on the more specific proposition should be modeled using independence but the effect on the more general proposition (GR) should not.

choose to make more than one person seem to see Stonehenge. And the same is true for many other aspects of the evidence. My visual experience of what appears to be an object leads me to expect to be able to have a corresponding tactile experience, and so forth. The tremendous power of GR to unify data as compared to GD is an important part of the confirmation of GD by the available data.

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 “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 of 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 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, whether human or non-human. 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.

Here--in the area of apparent predictability and perdurance--the disparity in likelihoods between GR and GD is particularly glaring. For while there is no problem at all accounting for these appearances on GR, GD gives us, again, a much greater degree of 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, 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 among them 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. Through 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 verisimilitude of his deception.

The term “verisimilitude” is itself relevant to the issue of likelihoods. As Lawrence BonJour (1999, p. 244) 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.” 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. This is particularly clear when it comes to perdurance and predictability.

Because of the fairly high likelihood of these aspects of our experience on GR, it is even easier to see the likelihood disparity with GD and the relevance of not making GD by definition empirically equivalent to GR. For without such a “piggy-backing” assumption, GD gives us no good explanation at all for the sheer quantity of data we have that seems to support a perduring and predictable physical world.

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. (Indeed, on any but a Berkeleyan version of GD, all the people you think you know or have ever heard of are themselves 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. For all of this, GD gives us little explanation at all except for the explanation we find buried in a sub-hypothesis of GD, a sub-hypothesis that we consider only because nothing else will even begin to do, and that sub-hypothesis simply tells us that for reasons unknown the deceiver has decided to do all of this, where “this” ends up being nothing other than a presentation to us of the evidence itself. Yet that sub-hypothesis is itself merely a tiny sliver of GD. Thus the evidence strongly and progressively confirms GR over GD.

GR and GD as natural rivals

Now that I have argued that GR is very heavily confirmed over GD by our evidence, it is clear that my initial decision to treat GR and GD as natural rivals--to treat them as if they have equal non-zero prior probabilities--is central to my argument. The more common approach to the

problem of the external world has been to place GR in rivalry with a sub-hypothesis of GD which is declared to be by definition empirically equivalent to GR and then to treat the entire issue as an interesting example of the problem of empirical equivalence. By treating GR and GD as rivals instead and by arguing that GR is very strongly confirmed by the evidence over GD, I have tacitly treated the prior probability of any sub-hypothesis of GD that might be empirically equivalent to GR (if such a sub-hypothesis exists) as much, much less than the prior probability of GR. How, then, do I justify treating GR and GD as rivals?

The point made earlier deserves to be repeated here: It is legitimate to treat GR and GD as natural rivals because GD appears on the face of it to be a classic deceiver hypothesis and hence to be a perfectly fair rival (fair to the skeptic, that is) to set against GR. GD appears different from the classic deceiver hypothesis only when one realizes that in setting it up I have not *defined* it as empirically equivalent to GR. But why should the skeptic be entitled to set not only the likelihood of a deceiver hypothesis *but also the prior probability* equal or nearly equal to that of realism? McGrew's original argument shows us that the prior probability of realism and of any empirically equivalent deceiver hypothesis are not equal in any event, so the question is simply how great the difference is. The skeptic might at that point (if he granted McGrew's argument) retreat to demanding that the difference between the priors be set to something insignificantly small, but there, too, we are entitled to raise this question: If we specify a deceiver hypothesis like GD that grants the existence of a deceiver who is to some extent both able and willing to make us believe that realism is true, why are we obligated to narrow our focus to a sub-hypothesis of GD which is empirically equivalent to GR for our present evidence and to make the prior probability of that sub-hypothesis nearly equal to that of GR? It seems that the only possible argument for such a move is precisely that by so doing we make it impossible for the evidence to tell us anything of much interest about the two hypotheses! This is hardly a rationale that should commend itself philosophically.

It is also worth considering that in treating GR and GD as natural rivals I have not "packed" GR so as to guarantee our evidence. Indeed, I have repeatedly acknowledged that much of our evidence will not have particularly high probability on GR. It is conceivable, for example, that there should be a real external world to which we have relatively reliable sensory access which is more boring, less detailed, and has fewer entities in it than the world in which we appear to find ourselves. That the evidence is so much more strongly unified by GR than by GD is a function of the actual evidential situation, not of an invidious attempt to make GR entail or even give extremely high probability to the evidence.

Another important consideration is that the treatment of GR and GD as natural rivals captures the intuition that there could be stronger or weaker evidence for GR over GD, which would not be the case if we pitted GR only against an hypothesis defined as empirically equivalent to it. The point is important and takes us back to the Case of the Clever Burglars. If the friend is allowed to declare his CIA-agent hypothesis to be by definition empirically equivalent to the hypothesis that nothing unusual happened the previous night, without specifying a set of evidence that we already have for which it is empirically equivalent, one consequence is that it is literally not meaningful to speak of further evidence for the normal night as better or worse vis a vis this rival. If I had, for example, scattered flour over the entire floor in front of all entrance points to the house, and if the flour were discovered apparently undisturbed the next morning, this would (given such a set-up) not provide any significantly stronger evidence against a break-in than if all doors and windows had been left unlocked with a clear

path in front of them. And it is easy to see why: The friend's hypothesis as defined can never be disconfirmed vis a vis the normal night hypothesis by *any* evidence; hence, his hypothesis cannot be *more disconfirmed* vis a vis the normal night hypothesis by one set of evidence than by another. If, on the other hand, we start with stable prior probabilities for clever CIA agent entry (in general terms) and for a normal night and progressively condition on the available evidence, we can see that more evidence that the house was undisturbed is (of course) better than less. The flattening out of all differences in evidential strength occurs only if we permit the paranoid friend to define his hypothesis in advance as empirically equivalent to the normal night hypothesis *no matter what the evidence* and by this means to avoid actually conditioning on the evidence that comes in.

Something similar is true in the deceiver case. It seems intuitively that we must have better evidence for realism and against the existence of a deceiver if we have many senses than if we have one, if we have clear sight showing us many details than if we have fuzzy vision, if we have multiple apparent reports than if we seem to be alone in the world, and if we seem to encounter many objects rather than few. But if we treat GR as the natural rival of a deceiver hypothesis that is *by definition* empirically equivalent to it, these differences cannot be reflected in our evaluation.

The problem of better or worse evidence is closely related to the issue of letting the evidence speak, which takes us back to the issue of *ad hocness*. The paradigmatically irrational *ad hoc* move is a refusal to allow the evidence to speak by refusing to admit disconfirmation. One particularly bold way of doing this is to declare when apparently disconfirmatory evidence begins to come in that, no matter what the evidence is, no disconfirmation can take place, because one's preferred hypothesis must be taken to be by definition equivalent to the hypothesis that would otherwise seem initially to be confirmed. In the case of L. Ron Hubbard and Dianetics, it would have been bold indeed for Hubbard to declare (as apparently he did not) that even if no protege of his were *ever* able to demonstrate perfect recall, he would always have an "explanation" handy and hence that Dianetics could never be disconfirmed by such failures. Had he done so, this would hardly have been a clever move on his part which would create an epistemic puzzle for philosophers throughout the ages (as the problem of realism and the deceiver is taken to do) but rather a blatant, and fatal, admission of epistemic malpractice.

Yet when we treat the problem of the external world as a pure exercise in empirical equivalence, we are enshrining the refusal to let the evidence speak, the paradigmatically irrational *ad hoc* move, in our very set-up of the case. It is not trivially true that deceiverism is empirically equivalent to realism, yet the usual treatment of the problem of the external world simply decrees the only deceiver hypothesis under consideration to be empirically equivalent to realism for any empirical evidence we might adduce now or later and then tells us (no surprise) that the evidence cannot distinguish between them! We avoid this consequence only by considering deceiverism in a form that could in principle be disconfirmed vis a vis realism and then asking ourselves what the evidence says about deceiverism--defined by its own *prima facie* general content but not by its relation to an unspecified set of possible evidence--and realism.

The Uninteresting Deceiver objection

At this point, the skeptic could attempt an objection that does not involve an insistence on treating realism and deceiverism as (by definition) empirically equivalent. The skeptic might

argue that by including deceivers who, like Melichus, cannot sustain their deceptions long-term or do not wish to do so, I am trivializing my result. After all, GD states that a *powerful* deceiver wishes to create the illusion of realism, yet in arguing for a very large likelihood disparity for the evidence, perhaps I have left this aspect of GD itself behind. How powerful, and how interesting, can a deceiver be if he would allow his apparent world to lapse?

To return to the in-world deceiver example of the CIA agent, the skeptic's complaint at this point might be similar to my friend's complaint if I were to argue that his CIA hypothesis is disconfirmed by the discovery that there are no fingerprints on his computer. *Of course* the CIA isn't going to leave fingerprints, he would point out. Even a garden variety criminal knows enough to wear gloves. Similarly, it is not much of an objection to his hypothesis to point out simply that I did not hear anyone in the night. It appears to be reasonable to expect that experienced, clever, and careful agents are capable of breaking into (at least an ordinary) house quietly enough not to wake anyone.

A related point is that I am giving zero likelihood for the evidence to (\sim GR & \sim GD). Even if we regard this as something of an idealization (that is, even if we admit that the likelihood for that middle hypothesis is merely very nearly zero), it seems that this aspect of the model requires placing within GD many deceiver scenarios that give the evidence, or at least some of the evidence, something noticeably higher than zero probability but that are still *very far* from being empirically equivalent to GR--deceivers who sustain the deception only for a few days, for example, or who are capable of or interested in producing only one type of sensation or only one other apparent human subject to interact with the deceived subject. The skeptic might argue that this inclusion of what he would consider to be "mostly incompetent" deceivers renders the conclusion of a large likelihood disparity between GR and GD uninteresting.

It certainly is difficult to specify exactly how competent and committed a deceiver or set of deceivers is included in GD. Terms like "powerful" and "wants to" admit of degrees. To make things more complicated, there are multiple axes along which a deceiver might be more or less competent and motivated; the multifaceted and multilayered nature of the evidence I have adduced for realism guarantees that this is so.

But I would dispute the contention that a deceiver who is able to produce only a small percentage of the evidence that we have is "uninteresting." Even Melichus is interesting. He is able to produce doors which, if they have no rooms behind them, nonetheless can be touched and can have a knife driven into them. His people are boring conversationalists and sometimes lapse into gibberish, but to be able to produce an inn full of pseudo-people who can make conversation at all is a good trick. Some of his pseudo-things can even be smelt. This is all very impressive despite the fact that he sustains the illusion only for a single evening. It would be still more impressive if he could do it for a full twenty-four hours.

To take another example, a deceiver who could produce the sensations and regularities experienced by a well-cared-for, healthy newborn infant would be an interesting deceiver. The newborn cannot get up and investigate his world with all his senses, as he will be able to few months later when he is able to reach and crawl, but he experiences regular satisfaction of his desires (apparently) in response to his cries. His visual and tactile experiences (as of the blanket, the light in the room, and so forth) as he moves his head and hands are regular and predictable, and, if we assume that he has a sensation as of time passing, this seems to go on for days at a time. The apparent visual experiences of his mother are associated predictably with a familiar smell and the familiar sounds of her voice. If a deceiver could do even this much, he would be

quite interesting even if he could not produce or did not choose to produce the richness, variety, and detail of sight, sound, smell, taste, and experimental results that the child will experience a few months later when he becomes mobile.

Moreover, the inclusion of a variety of types of evidence which realism unifies, so far from being a problem, provides multiple ways in which the existence of even very powerful deceivers (who were capable of providing all of one type of evidence but not others) is disconfirmed. So too do the various arguments for “overkill” in our evidence from the perspective of a deceiver. For example, I have pointed out that a deceiver might produce a great many regularities and apparent objects (no mean feat) while not giving us the apparent ability to discover laws explaining those regularities. The repeated use of the “overkill” argument shows that there is significant evidence for GR even against more powerful deceiver sub-hypotheses included in GD.

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 evidence--other, that is, than the evidence on which one is conditioning at the moment or for purposes of some given argument. From this perspective the distinction between the justification of prior probabilities and the conditioning 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 evidence into that which acts as background information and justifies the prior, on the one hand, and that on which one conditions 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 empirical propositions so sweeping and so vital as background to all of our other inferences about both everyday matters and high-level scientific theories that, prior to those inferences, there is no other empirical data to work with. This is paradigmatically the case when we come to consider the propositions contained in GR--the relative reliability of the senses and the existence and stability of the external world. If GR is strongly, even overwhelmingly justified, the problem of other minds will (plausibly) fall into place in short order, giving us, together with GR, a basis for all of our other empirical reasoning. But can we say anything definite about real prior probabilities of GR and GD? How are such priors to be justified in relation only to evidence *other than* the vast quantities of sensory and memorial data which we wish to use to find their posteriors?

A special problem arises here for the classical foundationalist, and as I am a classical foundationalist, I must confront that problem: The classical foundationalist holds that no proposition has an intermediate probability--a probability between zero and one--except in relation to some other evidence. There are no such things as propositions which are intrinsically probable without being intrinsically certain. Nor are there intrinsically improbable propositions. All intermediate probabilities are such in virtue of the relation of the proposition in question to some other propositions. 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 criteria for being foundational in such a scheme; the entire discussion here would be unnecessary if either of them did. If they have real, objective prior probabilities, then, those prior probabilities must be intermediate and, hence, must be explicable in terms of their relations to other propositions.

The problem here might, in theory, be solvable in terms of propositions knowable *a priori*. If we could divide up all possible states of affairs by means of symmetry conditions and then argue that GR and GD should have the same number of such “slices of reality” (or at least that GR should have no fewer than GD), we might generate intermediate probabilities like those used in our model here in relation only to tautologies. We might think of such prior probabilities as those knowable by a logically omniscient subject, self-aware but existing in a state without any sensory inputs and with memorial experiences only of his own previous thoughts. That is in theory. But whether or not it is possible in principle to do this, I have no story to tell about how such an approach actually does justify real equal prior probabilities for the two generic theories (or prior probabilities favorable to GR over GD). If we could ask the logically omniscient subject in such a state how probable it is that GD is true (or that GR is true) prior to all empirical evidence and relative only to tautologies, he might tell us that the question is meaningless.

It should be clear that the personalist avenue is not open to me, either. A strongly personalist approach would make the entire project of this paper pointless. If one is a personalist, one lacks the resources for condemning the paradigmatically *ad hoc* move, so long as the subject who engages in it always has coherent distributions. One might as well simply set one's prior probability for GR high to begin with--or, if one wishes to be a skeptic, for GD--and go on to do one's best with other, more practical inferences. No principled objection could be made to either move on a purely personalist approach.

If, then, I am not arguing *a priori* for real, specific prior probabilities for GR and GD, what is the value of the argument so far? Am I assuming something invidious in making that argument?

To answer that question it is important to stress that nothing in the model I have presented requires specific prior probabilities for either GR or GD, nor does it require that those probabilities fall within some range (e.g. “not very low”). I am arguing that GR is strongly justified in a way that can be seen even if both GR and GD are thought of as having very low prior probabilities. Because (\sim GR & \sim GD) has been assigned zero likelihood for our evidence (that is to say, the evidence we have is taken to rule out the negation of both GR and GD), and because I have assumed (uncontroversially) that both GR and GD have non-zero likelihood vis a vis the evidence, GR and GD could both “start out,” as it were, with very low prior probabilities and both end up much higher than they started. If nothing else, and aside from their relation to each other, both theories “take” probability from theories such as “this is a black hole universe with no knowing subjects” or “this is a universe in which we have nothing that appears at all like reliable sensory access to a real world.” The fact that neither specific prior probabilities nor high prior probabilities, nor even prior probabilities higher than some cut-off, are required for the inference should go at least some way towards mitigating the problem of the priors.

As mentioned above (see note 5), we can regard the model used here as simply normalizing the likelihood function for the evidence we have over GR, GD, and the negation of both GR and GD. 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 at all. *In effect*, this acts like a model in which GR and GD are given equal (but unspecified) priors with the prior of the negation of both being unimportant, since it is eliminated by the evidence. It seems to me, then, that the only epistemic question of any importance for the answer to the external-world skeptic is this: Is it fair or unfair to external-world skepticism to use a model that treats GR as if it is at least equal in probability prior to all empirical evidence to GD? Indeed, if the evidence for GR is as strong as I have argued it is, we might even ask further whether it is unfair to skepticism not to treat GD as having a *much* higher prior probability than GR *ab initio*. To object to this model, it seems that the skeptic would have to argue for the existence of some real absolute prior probability, relative to *a priori* truths for GD that is far higher than that of GR and hence is not overcome by the evidence adduced here, and that seems like a very difficult thing for the skeptic to do.

What, then, have I shown if the argument of this paper is successful? I have shown that the set of *all the empirical evidence we have* very strongly favors realism. I have shown that if we treat realism as being at least epistemically equal to deceiverism before any of the evidence comes in, the evidence overwhelmingly justifies realism. If we are concerned about following evidence, if we want to know where evidence leads, which way it points, and how strongly it points that way, this is no insignificant conclusion. It is difficult to see how such a conclusion could be trivialized by a difficulty in giving an *a priori* argument for real, specific prior probabilities--prior to all evidence--for these two hypotheses, especially when such specific prior probabilities are not required for the model in any event. Nor does the absence of such an *a priori* argument for real, absolute priors seem to give the skeptic any handle for an objection, since it is the skeptic himself who would need to make an argument for real priors justified *a priori* that so strongly favor deceiverism as not to be overwhelmed by the evidence.

If we were to ask our logically omniscient subject in a sensory vacuum what he thinks about the problem of the priors for this inference, this is what I think he might well say: "There are no specific priors in my epistemic circumstances for those two hypotheses, but the rational thing to do is to find out which one is justified over the other when you obtain some relevant evidence. So wait and see, and when you have some evidence, draw the conclusion it favors." If we follow that advice, we will not allow ourselves to be paralyzed by the problem of the priors. We can let the evidence speak.

Conclusion: A thought on direct realism

Despair of actually defending the propositions contained in GR on a more fundamental basis has formed 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 rational inference. Classical foundationalism is all too often associated with skepticism because of the classical foundationalist's interest in first philosophy--in giving an answer to the external-world skeptic rather than simply dismissing his questions out of hand.

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, tacit trust, trust one would be hard-pressed to defend explicitly, is incompatible with rational inference, then one will naturally have a problem with the attempt to justify these beliefs--still more to say that they must be justified--by way of inference, including a non-deductive, Bayesian inference such as the one presented here.

But what if we reject the opposition between spontaneity and tacitness, 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 fascinating possibility: These beliefs are so spontaneous, so pervasive, seem so obvious, and act as a background to so much else 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 so surprising upon reflection. Where a tacit inference is justified by so great a wealth of 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 to articulate his reasons for it. This need not be a sign of the absence of good reasons; it can be a sign of the presence of evidence satisfying one's reason to excess.⁸

This view of the matter also quite satisfactorily accounts for the fact that we would consider a person insane if he seriously did doubt the existence of the external world, which in turn accounts for the odium attached to the name of Descartes for pointing out that it is even possible to doubt it at all. 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 our own specific evidence. To all of these points the classical foundationalist can say this: Yes, you are right. Yes, there is a real 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, indeed, go on doing so. And let us go on doing so because what compels us in these beliefs is reason itself--reason which has been given so much material to work with, material which 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 sometimes 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 answering the questions we raise fully and thoroughly, so that the ordinary man need not doubt that, in the matters that are most important, the seemingly hyper-scrupulous and questioning reason of the philosopher is, ultimately, not opposed to but on the side of sanity, normalcy, and common sense. Thus we are both philosophers and real men living in the real world, and we see no conflict between the two.

⁸The difficulty for the ordinary subject in explaining his justification for belief in the external world is further accounted for by the fact that the clearest metalevel understanding of that justification comes through the use of probabilistic models and theorems (such as the theorem concerning the unification of data) that an untrained subject, and for that matter many highly educated people, will not have been taught.

Appendix: Berkeleyan Idealism

It is more difficult than it might at first appear to know whether to consider Berkeleyan idealism to be a version of GR or of GD. Berkeley himself would certainly have insisted that his view be considered a version of GR. This follows from Berkeley's belief, bolstered by his argument for *esse est percipi*, that the view of material things as the ideas of God (made available to finite spirits) is the only philosophically tenable view of matter and hence that it is knowable by philosophical reason alone. Berkeley's God as Berkeley conceives him is therefore not a deceiver of any sort at all, not simply because he is the Christian God and hence cannot deceive but, more importantly, because he does not hide his actions. We know, Berkeley thinks, if we just stop to think about it clearly, not only that he exists but also that what we call the external world must be literally a set of his ideas which he has made available to us.

One could, then, argue that Berkeleyan idealism should be considered a version of GR because, according to Berkeley, idealism just is the only tenable, correct understanding of what we usually call "the physical world" and also because God has chosen to give us relatively reliable access to these ideas of his.

It becomes difficult, however, to maintain that Berkeleyan idealism is a version of GR once one rejects Berkeley's Master Argument, as I do. At that point, the view that the external physical world is literally a set of God's ideas becomes merely another hypothesis about the causes of our experiences. It is no longer a philosophically knowable fact, knowable prior to investigation and evaluation of specific empirical evidence. Indeed, if a significant portion of our evidence for the existence of God, and especially of the Christian God, comes through empirical investigation, as I believe it does, then we must first believe GR before we can gather and evaluate empirical or historical evidence as evidence about something really happening outside our own minds. The same problem arises for a metaphysical argument (such as a cosmological argument) for the existence of God which starts with a simple proposition such as that the physical objects around us exist, since *without* the acceptance of Berkeley's argument for idealism, those premises will rationally be intended to refer to the existence of a real material world rather than to a world of Divine ideas. If Berkeley's strictly philosophical argument for idealism fails, then short of some later revelation from God himself that the physical world has been, all along, merely a set of his ideas to which he has chosen to give us access, it is difficult to see how we could ever know Berkeleyan idealism to be true.

This reflection suggests what I think is philosophically correct--that Berkeleyan idealism should be regarded as a version of GD. Even for a Christian, there is no particular reason to believe Berkeleyan idealism. For even if one believes that the Christian God exists, one does not have *ipso facto* anything like good reason to believe that the external world is literally the ideas of God. There is nothing to require such a view in the Bible or in traditional Christian teaching. And in fact, omnibenevolence argues strongly against Berkeleyanism (sans compelling philosophical argument), since without such an argument man is going to believe in a mind-independent physical world, which raises the specter of God as a genuine deceiver. (In other words, if there is no compelling and purely philosophical argument for Berkeleyanism, then an omnibenevolent Berkeleyan God is, plausibly, an incoherent concept.)

Something similar is true of the connection of the individual subject with other minds in a Berkeleyan universe. Berkeley is no solipsist. He posits the existence not only of God but also of multiple finite spirits. This might seem to be helpful to his theory, since his theory does not mean

that all our beliefs about other minds are false. But once we reject his own master argument for idealism, we are left with the uncomfortable problem of the status of other people's bodies. Our normal route to concluding that other persons exist is physically mediated communication with other beings who seem to us to be like ourselves. In Berkeley's universe, these people's bodies are themselves the ideas of God to which he has given us access, and our beliefs about the existence of other minds, like some of our arguments for the existence of God, must thus be mediated crucially through a false premise--namely, the existence of the mind-independent physical bodies of other persons, bodies which God has caused us to believe in by giving us sensory evidence but which, in fact, do not exist.

If we waive the problem of possible incoherence just raised and consider only likelihoods on a putatively coherent Berkeleyanism, it is true that, since Berkeley's God is omniscient and omnipotent, he does not have the potential failings of some of the possible deceivers we have considered. If an all-powerful God is the deceiver, he will never grow tired or bored, he will never lack power to effect his desires, and he will never lack the creativity to make his external world ideas detailed and convincingly real to us. This takes care of many of the likelihood problems we discussed above.

Even if Berkeleyan idealism did have equal likelihood to GR vis a vis the evidence, it should still by the above argument have a far lower posterior probability than GR. In that case, the radical disconfirmation of GD by empirical evidence should be regarded as a process in which the probability of GD continually converges on the probability of Berkeleyan idealism--that is, Berkeleyan idealism would be regarded as *the* one empirically equivalent deceiver hypothesis which is, in the end, compatible with our evidence while being much, much less probable than realism.

But I do not think that even this is correct. For even though God's power, knowledge, and creativity help the likelihood for our evidence to some extent, one can still imagine a Berkeleyan God who sustains our access to this particular set of ideas (ideas as-of-an-extramental-physical-world) for a time and then chooses to deal with man in some different way. God might, for example, show us this set of ideas for a time and then tear the veil away. He might allow finite human spirits to know one another by direct ESP rather than by the cumbersome process of what merely appears to be physically mediated communication. Since we have, on a Berkeleyan view, no material bodies anyway, doing so would not require destroying our bodies. (Indeed, it is difficult to see what "killing" a human being even means on a Berkeleyan view.) God could just suddenly stop giving us finite spirits access to those of his ideas which we call "our bodies" and "the bodies of our friends" and allow our finite minds to speak to one another directly. Or he might decide after a time to allow all men to have some more direct communication with himself without the mediation of seeming-matter. He might sometimes teach men by removing their perception of his external-world-like ideas and substituting some radically different ideas for those ideas.

Berkeley's God has not really engaged in an act of what would informally be thought to be creation at all--a point upon which Berkeley allows Hylas to challenge Philonous.⁹ Philonous's

⁹Philonous answers, "May we not understand [creation] to have been entirely in respect of finite spirits; so that things, with regard to us, may properly be said to begin their existence, or be created, when God decreed they should become perceptible to intelligent creatures, in that order and manner which He then established, and we now call the laws of nature?" (Fraser 1901, p.

response is that creation simply is God's choice to give some finite spirits access to some of his ideas, so that the creation in Genesis of the sun, moon, and plants prior to earthly sentient beings simply amounts to God's allowing other finite spirits such as angels access to his ideas of the sun, moon, and plants.¹⁰ But that sort of willed and purely mental access--as opposed to a true Divine decision to create a whole natural order of extramental entities and laws--seems remarkably arbitrary. It could be changed radically, as just discussed, without even the necessity of destroying the world. Since it is that notion of arbitrariness in the connections among our experiences that underlies part of the intuition of a higher likelihood of GR over GD, and since this arbitrariness remains even in Berkeleyan idealism, GR seems to have a higher likelihood for our evidence even than Berkeleyan idealism, though the latter is the best candidate for a deceiver hypothesis.

This analysis allows us to address the question as to whether GR is confirmed over GD by additional evidence of the same types we already have with or without upper bound. My central goal has been to argue that realism is strongly justified, not necessarily that its justification increases in principle with additional evidence without upper bound. If, for example, the rational probability of GR converges on something very high, such as (say) .9999, this means that the skeptical challenge can be answered. But the argument that even in Berkeleyan idealism, where we have eliminated problems of ability and are postulating a deceiver with limitless creativity, there is an ineliminable arbitrariness to the continuation of the deception suggests that the probability of realism may well increase with additional similar evidence without upper bound after all. While the argument of this paper does not turn crucially on that conclusion, it is nonetheless a highly satisfactory point in itself.

474)

¹⁰“In answer to that, I say, first, created beings might begin to exist in the mind of other created intelligences, beside men. You will not therefore be able to prove any contradiction between Moses and my notions, unless you first shew there was no other order of finite created spirits in being, before man.” *Dialogues* III, *

Works Cited

- Bellairs, John (1978) *The Face in the Frost*. New York, NY: Ace Publishing.
- BonJour, Lawrence (1999) "Foundationalism and the External World," *Noûs*, Vol. 33, Supplement: Philosophical Perspectives, 13, Epistemology, pp. 229-49.
- Fraser, A.C., ed. (1901) *The Works of George Berkeley D.D.*, vol. 1. Oxford: Clarendon Press.
- Gardner, Martin (1957) *Fads and Fallacies in the Name of Science*. Mineola, NY: 1957.
- Jeffreys, H. (1948) *Theory of Probability*. Oxford: Clarendon Press.
- McGrew, Timothy (1995) *The Foundations of Knowledge*. Lanham, MD: Littlefield Adams.
- _____ (2003) "Confirmation, Heuristics, and Explanatory Reasoning." *British Journal for the Philosophy of Science* 54 (4):553-567.
- McGrew, Timothy and McGrew, Lydia (2009) "The Argument from Miracles: A Cumulative Case for the Resurrection of Jesus of Nazareth," pp. 593-662 in *The Blackwell Companion to Natural Theology*, ed. W. L. Craig and J. P. Moreland. Oxford: Wiley-Blackwell.
- Rosencrantz, R.D. (1981) *Foundations and Applications of Inductive Probability*. Atascadero, CA: Ridgeview Publishing Company.