

YOU JUST BELIEVE THAT BECAUSE...¹

Roger White
MIT

1. Introduction

I believe that Tom is the proud father of a baby boy. Why do I think his child is a boy? A natural answer might be that I remember that his name is ‘Owen’ which is usually a boy’s name. Here I’ve given information that might be part of a causal explanation of my believing that Tom’s baby is a boy. I do have such a memory and it is largely what sustains my conviction. But I haven’t given you just any causally relevant information, I’ve given my *grounds* for my belief. I’ve given *reasons* that might *justify* me in supposing that Tom’s baby is a boy.

Less naturally, the question might be taken as a request for a broader causal explanation of my holding this belief. Appropriate answers might cite all manner of facts concerning the evolution of the human race, why I chose to pursue philosophy and hence came to know Tom, the mechanisms of email transmission, the firing of various neurons, the circumstances of concept formation as a result of which I’m able to grasp the thought that Tom’s baby is a boy, and so on.

It is an interesting question what distinguishes the narrower set of answers that I first suggested. I won’t pursue that here. I assume you have a good enough sense of the distinction I’m drawing. We might call the narrower set of answers *justifying reasons*, the kind of reasons I might cite in justifying my belief. Answers of the first sort are clearly relevant to epistemological evaluation. In assessing whether you know *p* or are rational in believing it to the degree you do, I will naturally want to consider what reasons you have for your belief. In deliberating myself about whether to believe *p*, in seeking an answer to the question of whether *p*, I will naturally consider what reasons or grounds I have to suppose that *p*. But what I want to focus on here is how explanations of the broader sort bear on such questions as whether to believe *p*. From a third-person perspective we can ask, ‘In assessing the epistemic status of S’s belief that *p*, what is the relevance of causal information that lies outside of the realm of justifying reasons?’ From the

first-person standpoint we can ask ‘In seeking to answer whether p , how should such causal information affect my deliberations?’

At first it might seem that such broader causal information could have little relevance if any. Like any belief my belief that p can be traced back to innumerable causes from far and wide. Most of these seem obviously irrelevant to the question of whether p . A wise person just considers the reasons for or against p and seeks to form his opinions to respect these. But this might be a misleading way to put the matter. The question is whether having formed a view as to whether p on the basis of reasons of the usual sort, information concerning the cause of your belief might constitute a further reason to change your opinion on the matter.

What sort of causal information might be thought to be relevant here? Let’s start with some examples. It has occurred to most of us that had we been brought up in a very different environment — a different religious/political/moral upbringing, say — we would have very different convictions. This can be unnerving, and for philosophers it can strike close to home. G. A. Cohen (2000) notes that Oxford graduates of his generation (including himself) tend to believe in the analytic/synthetic distinction, while Harvard graduates do not. Of course those on either side can argue vigorously for their own view, are aware of arguments on the other side, and can explain why they are not persuaded. But there is something unsettling about the thought that one’s convictions can be traced back to an arbitrary choice of where to go to school. In a similar way, perhaps it is to be expected that natural selection would produce organisms with a stronger sense of obligation toward their own children than to others. Sharon Street (2006) argues that we ought to be suspicious of this coincidence between the supposed mind-independent moral facts and the convictions of creatures produced by a process which it seems could hardly be sensitive to these moral facts.²

It is harder than we might expect to say what is distinctive about the kind of causal information that raises these epistemological worries. A first stab would be something like this: a factor F is thought to have made a causal contribution to my believing that p , but F appears to be irrelevant to the question of whether p .³ But this is clearly too broad. Any causal factor that I don’t list as part of my justification to believe p is one that I take to be irrelevant to the question of whether p (otherwise I would mention them in my defense of p). But only a fraction of the myriad of such factors even appear to raise a challenge to the status of my belief. Rather than try to characterize the nature of the factors I have in mind straight off, I’ll assume you get the general idea from examples. My purpose is to try to understand why the epistemological problem arises (or at least why we might be tempted to think that it does).⁴

Often epistemological problems turn out to be bigger than we thought. Doubts raised about certain judgments may generalize to a class of supposed knowledge, and to a broader class, and soon have us wondering if we have any reason to believe anything at all. Philosophers of a Moorean bent will think that the argument has gone off the rails at this point.⁵ Radical skepticism is too implausible. If premises aimed at a more limited target turn out to entail

that we are not justified in believing much of anything then this just shows that we are imposing implausibly strong constraints on justification. Without getting into the question of skepticism and Moorean methodology I too will be setting aside radical conclusions in order to focus on what if anything is distinctive and interesting about the problems raised by the apparent causes of our beliefs. Arguments for radical skepticism are common enough already. The interesting question is whether considerations of what led us to believe as we do can lead to doubts that can be quarantined — whether there is a stable position of doubt about those beliefs we were nurtured with which does not depend on or lead to doubt about most everything.

2. Debunking

The apparent effect of learning of certain causes of one's belief might naturally be described as *debunking*. But there two different ideas to be distinguished here. First, the causal information might be thought to *defeat* one's justification to believe that *p*. On this picture, upon graduation Cohen may well have been justified in believing in the analytic/synthetic distinction as his reasons were perfectly adequate. It is only upon discovering that he would have believed otherwise had he gone to Harvard that he can no longer rationally maintain his conviction. On this view it is his *learning* that his belief was caused in a certain way that prevents him from being justified in his belief, not the *fact* that it was caused in this way. (Indeed, if learning of certain causal factors can defeat one's justification, then arguably *justifiably believing* (perhaps mistakenly) that such factor were at work should have the same effect).

On the other understanding of debunking, the debunker takes the causal facts themselves to be incompatible with your being justified. On this picture Cohen does not *lose* his justification for his philosophical convictions upon learning of the influence of his Oxford education. This influence shows that he had no such justification to lose. The causal *information* has at most a second-order epistemic effect in helping him see that he was never justified in the first place. We can call the first kind of debunker the *undermining debunker*. He seeks to undermine your justification by revealing to you information about your causal predicament. The second is the *blocking debunker*, he argues that facts about your causal predicament block you from ever being justified, whether you realize it or not.

3. Distracting Issues: Reassessment and Disagreement

In order to isolate the phenomenon of interest here we should distinguish it from some others that are closely related. It is common for people to describe their loss of religious conviction as precipitated by the thought "I suppose I

just believe this because I was brought up this way. I could easily have believed something completely different had I been raised in a different home.” But plausibly what is going on here in many cases is just that a thought like this occasions the reassessment of their beliefs on more general grounds. The doubter may be led to consider what grounds he has for his previously unquestioned convictions and find them wanting. But reflection on the causes of belief need not really be playing any epistemological role here. Perhaps the doubter is right to think that he never had any good reason to believe as he did. But if so then it should make no difference if he were then to discover he was mistaken about the influence of his upbringing (perhaps he was really a stubborn child, and his parents’ best efforts to raise him in the faith had little actual effect even though he ended up believing for some other reason). We can set such cases aside as of little epistemological interest. One can be inspired to reassess one’s beliefs in all sorts of ways. The interesting question is whether the causal background of these beliefs can have epistemological relevance itself.

The second phenomenon to distinguished from our target is *disagreement*. Often our evidence for a causal hypothesis is a widespread correlation between variables. It is not enough that most people in circumstance C believe *p*. What clinches the case is that people in *other* circumstances do not believe *p*. So the evidence I have that environmental factors help explain my belief in *p* typically entails that it is highly controversial whether *p*. Now the fact that many people differ in their opinions may itself provide a powerful reason for doubt. This is the familiar problem of disagreement that has received so much attention lately.⁶ I don’t want to get into the disagreement debate too much here. Although the issues are connected my interest is in what epistemological significance (information concerning) causes of belief has *apart from* other factors such as disagreement. One reason to distinguish these issues is that our intuitions about cases may well be responding to factors other than the causal factors that concern us here. Perhaps we can try to isolate the matter with respect to disagreement in the following way. Cohen learns that opinion is divided among well-informed philosophers concerning the analytic/synthetic distinction. That may be significant grounds for doubt by itself. Suppose it turns out that while there is just as much disagreement on the matter as he thought, it is randomly distributed with respect to graduate education. He can no longer think that his choice of grad school was the determining factor for his convictions. But how much comfort is this? He still knows that about half the people who graduate from one grad school or another are mistaken about the analytic/synthetic distinction. Perhaps the thought that “I probably just believe as I do because I went to Oxford, not Harvard” adds an intuitive kick. But it is not clear how much.

Indeed it can be hard to see how it could add anything at all. Pretend that Oxford and Harvard are the only two graduate schools and I’m going to school in order to figure out whether *p*. We want to compare the following two cases.

Correlation: Equal numbers of students go to Oxford and Harvard. Those Oxford graduates believe p : Harvard grads believe $\sim p$.

No Correlation: Half of all philosophy graduates believe p , the rest believe $\sim p$. But opinions are randomly distributed with respect to grad schools.

How is the difference relevant? The crucial point seems to be that in Correlation I can identify *why* one philosopher believes and another rejects p . Of course all I can really do is gesture at such an explanation. It is not all that clear why philosophers should tend to believe as their professors do. There are crude hypotheses of course. Perhaps students were bullied into accepting the local orthodoxy. Or perhaps they uncritically looked up to their teachers as gurus. But these extreme stories are not realistic, and surely the worry that Cohen raises doesn't depend on anything of the sort. We may suppose (perhaps too charitably) that both schools exemplified all the virtues of graduate education that we could ask for. Students were exposed to all the relevant literature, they were encouraged to think rigorously, critically, and independently, and so on. Still, (surprisingly!) the Oxford grads tend to believe p but not the Harvard grads. I take it the worry that Cohen is raising still carries its intuitive force. But why should the worry be lessened in No Correlation? Perhaps in this case there is no single unified explanation for why one person accepts and another rejects p , so there is no easy way to predict what one will believe upon graduation. But it is hard to see how this could be make the situation any better. Whatever is going on, still only half of all the graduates are getting the matter right! If there is an important epistemic difference between the cases it apparently has to be that in No Correlation I am better placed to look back upon graduating and think 'I see now that I was fortunate enough to arrive at the truth'. But it is not clear why this move is any *less* plausible in Correlation. These preliminary considerations suggest that Cohen's information about the etiology of his belief is adding nothing over and above the problem of disagreement. But the matter deserves closer attention.

4. In Search of the Source of the Problem

It is much easier to get a feel for the worry presented by Cohen's case than it is to identify it or present it as a compelling argument. While Cohen's presentation of the issue is arresting, when he turns to a more careful analysis the results are surprisingly underwhelming. His most careful presentation of the argument relies on complex higher-order constraints on rational belief: If you have good reason to believe p rather than q you must be able to justifiably believe that your grounds for p are better than another's are for q . And you can rationally believe p only if you don't know that you lack good reason to believe p rather than q . There is a certain plausibility to these conditions, but the following premise is given little support:

(3) In a wide range of cases of nurtured belief, people *who continue to believe p* (can readily be brought to) realize that they believe *p* rather than *q* not because they have grounds for believing *p* that are better than the grounds for believing *q* that others have, but because they were induced to believe *p* without being supplied with such differentiating grounds. (p. 13)

Cohen's main illustration involves twin sisters separated at birth. One is raised as a Presbyterian the other a Roman Catholic. They retain their respective faiths and meet later to reflect on the fact that had they swapped homes at birth they probably would have had each other's religious beliefs also. I can see one way that we could continue the story which would fit with Cohen's premise (3). The Presbyterian might think back to her childhood, "The distinctive doctrines that separate Presbyterians and Catholics were simply taken for granted when I was growing up. No one ever so much as addressed the question of why one should be a Presbyterian rather than a Catholic. And now I think about it, I can't think of any good reason myself." But this is just the less interesting case that we set aside earlier. Reflections on her upbringing are simply the occasion for a reassessment of her beliefs. In any event, it is not the kind of case Cohen has in mind. Cohen imagines the sisters vigorously arguing points of doctrine to no avail. They each present arguments with which the other was familiar and to which she has a considered response, and so on. If they each have so many dialectical maneuvers at their disposal it is hard to see why they would think that they had not been "supplied with such differentiating grounds."

Of course as the sisters argue on without resolution it will be natural for each to think, "If these arguments that I find so convincing really are rationally compelling, shouldn't they move a reasonable person to conviction? But my sister seems reasonable enough. Perhaps my reasons aren't all that I take them to be." This is just the problem of disagreement. But while Cohen doesn't carefully distinguish the issues he does make it clear that the worry he is raising is not meant to hinge on this. Cohen insists that the story could be told in our own case with a merely *hypothetical* twin with different religious or political views. But it is hard to see what threat a merely hypothetical disputant poses. When Cohen graduated he surely thought that he had plenty of grounds for his belief in the analytic/synthetic distinction. If we had asked him why he believed it he may have provided paradigm examples, appealed to a certain theory of meaning, shown what he took to be the flaws in opposing arguments, and so on. He surely thought that taken together these considerations supported his view over the view that there is no such distinction. He couldn't have become convinced in the first place if the arguments didn't strike him in this way. Now we suppose he imagines a *hypothetical* philosopher who thinks otherwise while being aware of all the same considerations. If he really thinks that his arguments are rationally compelling he should just think that any such possible philosopher would have to be less than fully rational in his judgment on this point. If he imagines that this hypothetical philosopher is fully rational in his judgment on the matter he

will have to think that the philosopher agrees with him.⁷ The questions “Which conclusion is supported by this evidence?” and “What would a rational agent conclude on the basis of this evidence?” are very closely related. It is *actual* disagreeing philosophers that are more stubborn. If he thinks his arguments are rationally compelling he will be surprised to find that apparently rational philosophers see it differently. We can put the difference in terms of explanation. It is a *necessary* truth that it is *possible* for someone to disagree with Cohen. No hypotheses concerning the merits of his arguments or the truth of his conclusion are needed to explain this necessary fact. That someone *actually* disagrees is a contingent matter, and one that is more to be expected given that his arguments are not so good than that they are. The force of the worry at least as Cohen sets it up here appears to derive solely from the phenomenon of disagreement rather than anything about his philosophical nurturing.

4.1. *Inevitability*

Let’s consider some other ways that we might try to identify the distinctive epistemological problem about the causes of one’s belief. One curious feature of our naïve thoughts about these matters is that there are different ways of further motivating skeptical worries from causal information that seem diametrically opposed. On the one hand some will emphasize the *contingency* or *chanciness* of the situation. It can seem to be at best a matter of *luck* whether I went to a grad school that would inculcate me with true beliefs. I could easily have had a different religious upbringing in which case I would most likely have a radically different view of the world. The precariousness of my having come to the truth (if I have) can be worrying. On the other hand, some will want to emphasize the *inevitability* or *predictability* of my believing as I do given my circumstances. “You were *bound* to believe that, given your background.” “That’s just what I’d *expect* someone who studied under Prof. X to think.” An extreme case might be someone who had come under the influence of a cult. “These people use methods that are very effective at getting people to believe whatever they want,” we might tell him, hoping that this will prompt him into doubt. Of course there is no inconsistency between these two factors, they just occur at different stages. It is a contingent matter which causes are in place that make it inevitable that you will believe this or that. But if both factors by *themselves* raise an epistemological problem then it would seem to be a matter of damned if you do and damned if you don’t. And that is reason to suspect that our naïve thinking on this matter is confused.

It is the emphasis on inevitability that I want to begin with. What is worrisome about the thought that my beliefs are the inevitable product of certain conditions? One thought is similar to familiar worries about free will and responsibility for action. Perhaps if there are forces determining that I will end up believing *p* then I don’t have the right kind of epistemic autonomy,

the ability to seek the truth and follow the evidence wherever it may lead. I'm just led along in my thinking by forces outside of my control. I'm not sure that there is much to this thought. Action and belief are importantly different. I have little voluntary control over my beliefs. Indeed those conclusions that seem *inescapable* are usually the ones that are so obviously rational. I couldn't stop thinking that $1 + 1 = 2$ if I tried, but that casts no doubt on it. My views about politics, sport, and religion are perhaps more malleable, but they are also more questionable. There are aspects of inquiry that are more under my control and questions of responsibility may arise here. Perhaps if I'm forced to do my arithmetic homework against my will I get less credit than if motivated from within by a passion for the truth. But this has no bearing on whether the conclusions I reach are *epistemically* justified.

Clearly causal determination of belief is not worrisome as such. I suppose it was all but inevitable that I would believe that there is a computer before me given that there is one and I'm looking at it. This kind of causal determination is just what we want for our beliefs. The same holds for other matters that don't necessarily involve perception. It was predictable that Cohen would graduate from Oxford accepting the importance of *quantifier-scope* distinctions in addition to the analytic/synthetic distinction. But concerning the former at least, Cohen should surely think "Well there *is* an important distinction between wide and narrow scope, and a good education is bound to reveal it." (Again, it is significant here that he knows of no credible philosopher who denies this). The problematic kind of causal influence would seem to be that which operates independently of the *truth* of our beliefs. The worrying thought is that given the forces acting upon us, we would have believed that *P* *whether it was true or not*. For instance, if it is evolutionary forces that have shaped our moral convictions and these forces are sensitive only to matters of trait fitness, not moral truth, then it would seem that we are bound to end up with our actual convictions regardless of the moral facts. So the problematic fact I might discover about my belief is perhaps the following

Insensitivity: If *p* were false I would still believe *p*.

Recall the two ways that debunking may be understood. The *blocking* debunker claims that a *fact* such as Insensitivity is incompatible with my being justified in believing *p* (regardless of my awareness of this fact). So on discovering that Insensitivity holds I do not lose my justification for *p* but discover that I was not justified to begin with. This version of the debunking strategy is a non-starter if the fact appealed to is Insensitivity. Any false belief of mine automatically meets the Insensitivity condition (*p* is false and I *do* believe it). But it is not the case that every false belief is unjustified.

The only promising debunking strategy then is one that takes evidence of Insensitivity as a *defeater* of one's justification. The thought must be that even if I am initially justified in believing *p*, I lose this justification once I have reason to believe that this belief is Insensitive. This suggests a general principle:

Truth Sensitivity: If I'm justified in believing that I would have believed p even if it had been false, then I'm not justified in believing p .

This principle can seem very attractive. If I have reason to think that I am not connected to the facts in a manner preventing me from error, then I have reason to doubt that I have in fact avoided error, and hence reason to doubt p . But the principle is almost certainly false as it stands. For it quickly leads to the conclusion that I'm not justified in believing much of anything. Suppose I'm justified in believing in at least one proposition p . Then I am clearly justified in believing the disjunction *either p or I don't believe that p* . Since this trivially follows from p I can't coherently be confident that p while remaining uncertain or denying the disjunction. For this disjunction to be false is for p to be false and yet I believe p . But in any case in which I believe p (mistakenly or not) I naturally believe the obvious consequence that either p or I don't believe it. So in other words if *either p or I don't believe that p* were false, I would still believe it (and of course I'm justified in believing that this is my predicament). Hence it follows by Truth Sensitivity that I can't be justified in believing that either p or I don't believe that p . But as we have already noted, I must be justified in believing this disjunction if I'm justified in believing p . So we have reached the conclusion that I'm not justified in believing p . That is, I'm not justified in believing anything. This is not the conclusion that the debunker is usually aiming for.⁸

This example has an air of trickery about it. So it might be worth noting that the problem arises more generally from cases of empirical underdetermination. As jury members we have overwhelming evidence that the defendant is guilty (finger prints, eyewitness reports, DNA matching blood on the murder weapon. . .). Although we are more than justified in believing that he is guilty, our evidence doesn't strictly entail it. There is always logical room for some wildly *ad hoc* conspiracy theory according to which the defendant was framed and all the evidence was planted. We will have to be justified in denying this conspiracy theory if we're justified in believing the defendant to be guilty (since they can't both be true). But if the conspiracy theory were true we would have all the same evidence and hence quite reasonably (but mistakenly) believe the defendant to be guilty and hence the conspiracy theory to be false. So again, by Truth Sensitivity we derive the absurd conclusion that we aren't justified in denying the conspiracy theory or believing that the defendant is guilty. What we have here is of course a familiar skeptical strategy: the denials of skeptical hypotheses ('I'm a brain in a vat', 'The sun has risen every day, but won't tomorrow' 'the world came into existence five minutes ago with false memories in place'. . .) typically fail a condition like Truth Sensitivity. And this leads quickly to radical skepticism. We could try to tweak the condition to avoid radical skepticism while retaining its skeptical consequences in a more limited domain like philosophy or religion. But I doubt we'll make much progress at that without it looking like a lot of *ad hoc* maneuvers.

Even if we suppose that something like Truth Sensitivity is correct, a further difficulty concerns the conditions in which I should *think* that my belief fails to

be counterfactually sensitive to the truth. A case that presents peculiar difficulties is theism. Suppose I'm justified in believing in the existence of God (prior to considerations of Insensitivity). Naturally I think that God is ultimately responsible for my theistic belief. Indeed I'm justified in thinking this if I'm justified in believing that there is a God. So of course I should think that my theistic belief is not at all independent of whether there is a God. If there weren't I wouldn't be here, let alone mistakenly believe in God. How could empirical investigation reveal to me that this is not so? It won't be enough to trace the origins of theistic belief to some prior physical condition. Perhaps we find that theistic belief is a byproduct of some cognitive trait that gave our ancestors some fitness advantage having nothing to do with the truth of theistic belief. Whatever causal condition *C* that we might trace theistic belief back to in our cultural or evolutionary past, it is only appropriate for me to think that this condition would not have obtained unless God had brought it about. How could I have reason think otherwise? Perhaps I could somehow obtain evidence that condition *C* was not in fact a part of God's plan at all (God might tell me as much). This sort of information probably won't come to light by ordinary scientific investigation, and in any event could hardly render my theistic belief unjustified. I might instead gain evidence that there is no God and hence that my theistic belief is counterfactually insensitive to God's existence. Then sure enough my justification may be defeated. But it will have nothing to do with learning of causal factors operating independently of the truth of my belief.

Similar considerations apply to the closely related issue of *explaining* religious belief. We hear it said that science is beginning to explain (or perhaps "explain away") religious belief, and that this somehow poses a challenge to the believer. It is not at all clear how this is supposed to work. Of course explaining a belief poses no threat to the belief as such. The explanation of my believing that *p* is often that *p*. Sometimes it is emphasized that scientists are beginning to explain religious belief without appeal to the *truth* of those beliefs. And this lack of connection to the truth ought to lead one to doubt their truth. But the truth of an explanation of my belief that *p* that makes no reference to whether *p* doesn't by itself pose any threat to the justification of my belief. I believe that David Hume was born in 1711 because I read in the introduction to my copy of the *Treatise*. Perfectly good explanation. No mention here of whether he *was* born in 1711. But that hardly disturbs my belief. To be even potentially challenging, it seems the theory would have to suggest that the *ultimate* explanation of my belief makes no reference to Hume's actual birth. The principle here would be something like

Explanatory Relevance: If I'm justified in believing that the ultimate explanation of my belief that *p* makes no reference to whether *p*, then my belief that *p* is not justified.

This principle faces a number of problems. I'm justified in thinking the sun will rise tomorrow. But assuming it does, this fact plays no role in explaining why

I thought that it would. It is hard to see how necessary truths, especially of a philosophical sort can explain contingent facts like why I believe as I do. For instance it is mysterious how Explanatory Relevance itself could play any causal role in bringing about my belief in the principle.

So Explanatory Relevance will have to be restricted in some non-*ad hoc* way for it to be of any use to the debunker. But even if we suppose that the principle holds at least in the specific case of theism, there is still a puzzle of how a theist could reasonably come to think that the ultimate explanation of his belief does not involve God. Of course it is not enough that we come up with some story which, if true, would explain it. That there's a typo in the relevant passage in my copy of the *Treatise* would, if true, explain why I think Hume was born in 1711. But such a possibility hardly undermines my belief about Hume's birth. We would need some reason to suppose that the explanation is correct. How might a theist come rationally to think that the correct ultimate explanation of his belief does not involve God? It would seem that for whatever cause *C* of theistic belief that science might reveal, a theist should just think that God is responsible for *C*, that that is just the process that God chose to instill belief in his creatures. It's not just that I *can* consistently suppose this in some desperate attempt to preserve my faith. It appears to be the only reasonable thing to think, assuming that my theistic belief is justified to begin with. Perhaps it could be argued that certain causal routes to belief are unlikely options for God, and so even a theist ought to be dubious that that was what God was up to. But I'm not sure how we would figure this out. And the would-be debunker of theistic belief would do well to avoid speculative theology.⁹

None of this should suggest that theism gets to enjoy some special status of immunity to epistemic defeat. One way to support the claim that the ultimate explanation of theistic belief does not involve God is to argue directly against the existence of God. Of course such arguments might have merit, but they are not my interest here as the epistemic defeat involved needn't have anything essentially to do with the causes of belief.

There is however a more subtle variant of this kind of objection.¹⁰

Suppose we make great strides in explaining religious belief in naturalistic terms. We will have gained a great deal of understanding about how such belief arises without yet supposing that there is a God. Now of course one can consistently tack on the further claim that there is a God behind it all. But it is *simpler* not to. So considerations of parsimony will favor atheism and hence potentially defeat one's justification for theistic belief.

Perhaps parsimony considerations count against theism quite apart from explanations for belief. It's not clear what to make of this. It's not as though the argument schema "a theory without Fs is more parsimonious than one with Fs. Therefore there are probably no Fs" is particularly compelling in general. But even granting that an objection from parsimony may have some merit, it still doesn't seem that the explanation of *belief* is playing a crucial role here. Instead

of theistic belief we might have picked another phenomenon, for instance that *grass is green*. We've learnt that grass is green because it contains chlorophyll. And a suppose we have a further explanation of why it does, and so on. The textbook gives a very thorough and compelling account of why grass is the color it is, but nowhere is *God* mentioned. As before, we can consistently tack on the claim that God is behind it all, but it's simpler not to.

Now I've heard a number of objections to belief in God over the years but I've yet to hear it suggested that scientific investigations into the *color of grass* pose a special threat to such a belief. Why then should the argument above carry any more force? Curiously, there are some explanatory endeavors in science that are thought to pose a special threat to theistic belief. The theory of natural selection is the obvious example. Understanding why this might be so may help us see why the argument above can be tempting. Many people have thought something like the following.

Prior to Darwin there was a quite compelling argument from design, a kind of inference to the best explanation. William Paley (2006) was right to think that the intricate machinery of the mammalian eye and other biological wonders cried out for explanation, and that the only compelling explanation available at the time was theism, or at least some kind of powerful intelligent agency. But Darwin (1998) changed that. In giving a *better* explanation (or at least the framework for such an explanation) Darwin undermined Paley's argument, thus removing this potential source of justification for theism.

In any event, I think something like the explanation above is the only credible way that evolutionary biology *could* be thought to have epistemological relevance to theism.¹¹ And it might seem natural to carry this thought over to the case of explaining theistic belief. But now for the story above to have the consequence that confirmation of the theory of natural selection would render one's theistic belief unjustified, two controversial assumptions must be made. First, that the pre-Darwinian argument from design was indeed rationally compelling in the first place. And second that this is the *only* source of one's justification for theistic belief, that one's justification crucially hinges on the success of this particular theistic argument. Whatever modest plausibility the conjunction of these claims may have, it is entirely lost when we shift from explaining the *eye* to explaining theistic belief. First, I doubt the fact that people believe in God strikes anyone as the kind of startling fact crying out for explanation in terms of design the way that biological marvels like the human eye can. Nor has anyone thought (unless they *independently* believe in God) that the hand of God is the only satisfying explanation available for the phenomenon of theistic belief. And anyone so generous as to find the Argument from Theistic Belief compelling grounds for belief in God is unlikely to think that this is the *only* source of justification for his belief. If one did hold this peculiar combination of views then I think we can see how parsimonious naturalistic explanations of theistic belief might

be thought to render such belief unjustified. But I don't imagine this actually applies to anyone.

I suspect that more generally what is really behind the appeal to Explanatory Relevance or the Truth Sensitivity condition is something like the following. In many cases where our justification for a belief crucially depends on a certain item of evidence that justification can be defeated by providing an alternative explanation of that evidence and so "screening off" its relevance. That the gas gauge reads Full is evidence that the tank is full. But now I learn that the gauge is stuck and would indicate Full whether the tank is full or not. In this way my belief formation was not sensitive to the truth as the gas gauge on which I based my belief was not sensitive to the contents of the tank. This information undercuts the evidential support of the gauge reading. And if my justification for believing that the tank is full crucially depended on this evidence then I am no longer justified in supposing that the tank is full. (Plausibly this only hinges on my being *justified in believing* that the gauge is stuck, whether it really is or not). A natural way to think of this is in terms of explanation. In the absence of an alternative explanation, that the gauge reads Full is best explained by the fullness of the tank. Once I know that the gauge is stuck I can explain the gauge reading without any appeal to the state of the tank. So we no longer have explanatory support for the hypothesis that the tank is full.

It is cases of this sort that likely lie behind the plausibility of something like Truth Sensitivity or Explanatory Relevance. But the lessons of the gas gauge case don't carry over well to the case of explaining belief. It is not as though I typically take that fact that I believe something as crucial evidence for its truth. I may do so with regard to the beliefs of *others*. Perhaps my only reason to suppose that your middle name is Rufus is that you told me so. If for some peculiar reason I come to see that you would think your name was Rufus whether it was or not then my justification for supposing that it is has been undermined (you are 'stuck' on the idea that your middle name is Rufus just as the gauge was stuck). But that's not the way that I think of my own beliefs. I don't think over whether *p* and then upon coming to the conclusion that *p* think '*p*', and now I believe that *p*. Smart fellow that I am it is unlikely that I would believe that *p* if it wasn't true. So this is *further* evidence that *p*.' and thereby increase my confidence that *p*. If it were legitimate for me to reason in this way and I did so, then perhaps giving an alternative explanation of my believing that *p* could undermine this extra boost I get from the consideration that I believe that *p*. This would be parallel to the explanation of the gas gauge. But this would still make little difference to the status of my belief.

4.2. Evolutionary Explanations

Evolutionary explanations of belief are a special case that is worth considering separately for a moment. Various stories have been told about how certain

religious or moral beliefs might have served some evolutionary purpose, or that dispositions to believe are a byproduct of other cognitive capacities that can give an organism a selective advantage.¹² There are different ways that one might think that evolutionary scenarios have a debunking effect. The first follows the theme of inevitability. Crudely put, natural selection is likely to favor creatures with certain moral convictions such as having an obligation to care for one's own family. This prompts the thought that we really only hold these moral views because evolution planted them in us. And this can be unsettling, for what kind of moral guide could blind natural selection be?

To look more carefully at what is going on here we can consider a pared down analogy.

Adam's Party: Adam throws a party and we're all invited. As we arrive Adam asks each of us whether *p*. You answer that *p* and go in to enjoy the party. We discover later that he had a gun in his pocket and was prepared to shoot anyone that didn't believe that *p*.

We don't actually know to what extent he used the gun. Perhaps all the guests happened to believe that *p*, or perhaps there were ten times as many guests invited but he has piles of dead bodies in his basement. We just know that none of us would be there if we hadn't believed that *p*. So the belief that *p* has a major selective advantage in this environment.

The first thing to note is that Adam's crazy plan does not make it inevitable, or indeed does anything to explain the fact that *you* happen to believe *p*. You came in believing *p*, and would have believed it regardless of what Adam was up to. What is explained is why *everyone at the party* believes *p*, i.e. why there are no *p*-doubters at the party (if there were any there to begin with they were selected against). So the thought that "I really only believe *p* because of selective pressures at work" is misguided. But aside from this, it is clear that when you learn of the selective advantage of your belief that *p* in this environment this has little or no relevance to the epistemic status of your belief. You may or may not have been justified in believing *p* in the first place. If you weren't then you are still not, and nothing has changed. If you were justified to begin with then Adam's plan of doxastic genocide does nothing by itself to rebut or undercut or call into question in any way that justification. Real cases of natural selection at work are rather more complicated than our toy case. They involve small changes conferring small advantages accumulating over many generations. But if there is no real epistemological challenge arising in our pared down story it would be surprising if it somehow arises from the intricate details of actual natural selection. It is not easy to see where in the details the epistemological relevance could be located.

We could add to the case some assumptions that make the selection of *p*-believers relevant to the question of whether *p*. Perhaps we think it likely that Adam might know whether *p* (although we don't know what his view is) and

that he really dislikes people with true beliefs (excluding himself). With this background in place Adam's shooting plan might provide some reason to doubt that p . Perhaps the only reason we have survived is that we were let off for being mistaken. Conversely, if we thought Adam despised false believers (or despised true believers but is himself anti-reliable with respect to P) then we would have gained some evidence for p . But we get closer to our actual evolutionary situation if we assume that we have no reason to suppose that his shooting policy has any significant correlation with truth or falsity on the matter. Or even if it does, we have no grounds for supposing that the correlation goes one way rather than another. This is more like the kind of evolutionary scenario that people have in mind with debunking arguments. For some cognitive faculties such as *vision* we might expect natural selection to favor a system that produces true beliefs since true beliefs about one's environment are useful for survival. But for others like our moral faculties or religious beliefs, natural selection would appear to be *blind* with respect to facts in these areas.

There is one minor way in which Adam's shooting policy might be relevant when assessing one's belief that p . Perhaps before learning of Adam's policy I ask others at the party if they answered that p . It turns out that everyone there believes that p . This might bolster my conviction that p to some degree. For I might think that if p is false it is likely that one of you smart folks would have found some reason to doubt or deny it. That you all believe it too is further evidence of its truth. The information that Adam only allowed p -believers to live screens off the relevance of this consensus. It was to be expected that I would find no dissenters at the party regardless of what reasons there are to believe or disbelieve that p that I might not be aware of. The result of learning that of Adam's policy in this case is just to put me back to the state I was in prior to learning of other partygoer's opinions. But I might remain quite convinced that p and justifiably so. This degree of possible relevance is appears to be fairly minor and in any event hinges solely on matters of agreement/disagreement which are not our concern here.

I have been focusing on one central factor in evolution, namely *selection*. Of course the other important mechanism at work is *random mutation*. I suspect that if there is an epistemological worry arising from evolutionary history it stems not primarily from considerations of fitness and selection but rather the thought that one's cognitive faculties are ultimately the product of a long series of random mutations among one's ancestors. This leads us to the a different way of thinking of the epistemological significance of causes of belief, which focuses on *contingency* or *luckiness*. I'll be looking at that more generally in the following section, but let's continue for now considering evolutionary scenarios. The worry would go something like this. What are the odds that a series of genetic mutations will result in creature with mostly true moral beliefs? Surely not very high. The initial worry raised here has nothing essentially to do with natural selection. Imagine a counterfactual scenario in which the environment is so cushy that organisms rarely have any difficulty reproducing.

The result is a plethora of organisms with all manner of traits, traits that would be a significant disadvantage in our actual environment. The variety of features randomly produced might include tendencies to form the whole range of possible moral judgments. That natural selection plays no role in this scenario in weeding out agents with different moral inclinations does nothing to alleviate the skeptical worry posed here. Perhaps it is likely, given a large enough number of mutated organisms, that there be *some* creatures who tend to make correct moral judgments (I have no idea how we could begin to calculate this). It remains as unlikely as ever that *I*, or my like minded species-mates, should win the evolutionary lottery and be blessed with reliable moral faculties.

It is here that natural selection can be relevant. It is not that considerations of selection pressure *create* the skeptical worry; they merely *fail to solve it*. With regard to some other cognitive capacities the situation may seem brighter. Concerning induction Quine (1969) suggested that “there is some encouragement in Darwin” (p. 13). As he nicely pointed out, “Creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind.” (p. 13) While it is unlikely that a particular sequence of mutations will produce reliable inductive reasoners, perhaps it is to be expected that those (if any) that survive over generations will be good at predicting what is going on about them. Not so for moral beliefs. There is no reason to suppose that getting the moral facts right gives one a better shot at reproducing. We could think of the comparison as follows. Suppose the Oracle tells us that on Planet X, evolution by unguided mutation and natural selection has been going on for some time and has yielded among other things some creatures capable of forming beliefs about their physical environment including local events in the short term future, as well as moral beliefs what they ought to do, what is a reason for doing what, and so forth. Without being told anything about the *content* of their beliefs, what are the odds that their beliefs are mostly correct? Considerations of selection pressures may give us higher expectations for the truth of their non-moral beliefs that have more obvious survival value. But that is just like *our* planet. Shouldn’t we be similarly dubious of our moral convictions more than perceptual beliefs and inductive conclusions?

More needs to be said to turn this into a rigorous argument for moral skepticism. In considering Planet X with just abstract evolutionary considerations to go on we could hardly predict that among the creatures there are some who play the bassoon, or ride tricycles, or study organic chemistry. But plainly on our planet there are. Why can’t we respond similarly to the debunker? We can admit that the minimal details of the process of evolution give us no grounds to expect there to be reliable moral thinkers. But we can insist that it is clear that that is in fact how evolution has turned out: On a broad range of fundamental moral questions our opinions are largely correct. The debunker will have to claim that there is a special problem with this response having to do with our reliance on our own moral judgments in order to judge whether evolution has produced reliable moral thinkers in our own case. She will likewise have to rule out a more

subtle response. We might deny that it is so unpredictable that evolution should produce creatures with correct moral beliefs. Ironically we can appeal the very kind of fact that debunkers emphasize. For example natural selection is likely to favor creatures with a sense of obligation toward their offspring. But necessarily, agents *do* have an obligation toward their own offspring. So it is to be expected that evolution will produce creatures with correct moral beliefs on this and a range of other matters. Again the debunker will have to rule out this sort of appeal to one's own moral judgments in assessing the likelihood of the outcome of evolution with respect to the reliability of these judgments. The issue here is crucial the skeptical argument yet receives no attention from Street or Joyce. Street illustrates the epistemological problem with a striking story

[On a realist theory of value that posits no relation between evolutionary forces and evaluative truth] allowing our evaluative judgments to be shaped by evolutionary influences is analogous to setting out for Bermuda and letting the course of your boat be determined by the wind and tides: just as the push of the wind and tides on your boat has nothing to do with where you want to go, so the historical push of natural selection on the content of our evaluative judgments has nothing to do with evaluative truth. (p. 117)

Street notes that by some extraordinary stroke of luck we *might* stumble across Bermuda, or evaluative truth, but suggests that we are left with “the implausible skeptical conclusion that our evaluative judgments are in all likelihood mostly off track.” (p. 117)

But the story here doesn't actually suggest a skeptical conclusion. Before we reach land we should judge it very unlikely that we will get where we want to go. (If we could accompany evolution on its path through genetic space, the prospects of arriving at evaluative truth should seem bleak). But once we come ashore and see the people and street signs and resorts there can be little doubt that we have arrived at Bermuda (even if we must marvel at our good fortune!). Why can't we similarly recognize evaluative truth when we stumble upon it even by extraordinary accident? That we have an obligation to care for our children for instance seems about as easy to recognize as anything.¹³ To make the nautical story more conducive to a skeptical conclusion we could suppose that the North Atlantic is littered with “Fake Bermudas”, large papier mâché facsimiles of Bermuda crafted to look and feel just like the real islands. If we are aware of the proliferation of these fakes then even if we are lucky enough to land in Bermuda we have reason to remain dubious. There is certainly room to argue that the evolution of evaluative judgment is relevantly like sailing among fake Bermudas. (The many possible wildly false moral theories will *seem* true from the *inside*). But these are substantive questions and I think this is where most of the epistemological action is.

The subtle matters arising here will be addressed more thoroughly in the next section. For now I want to consider whether the skeptical problem raised

here, if there is one, is really a special problem arising from what we know about our evolutionary history, or is really trading on ideas that lead once again to far more broad and radical skepticism. We can begin by comparing our situation to one of initial *agnosticism* concerning the origin of our cognitive faculties. Suppose we were entirely unaware of the theory of natural selection and have no idea how it is that we got to be here. Here we are finding ourselves irresistibly judging certain actions right or wrong. If a skeptic poses the challenge, “What grounds do you have to suppose that you are at all reliable at judging moral matters?” we will have little to say. The best we can try is to engage in moral reasoning to try to persuade her of the truth of our opinions (this will do about as much good as Moore raising his hands). If the skeptical worry beginning from knowledge of our evolutionary origins is on the right track, then there is surely no less of a worry arising from *ignorance* about our origins. If we had no clue as to how we came about how could we possibly predict (without appealing to our moral judgments) that our moral judgments would be at all reliable? This suggests that the evolutionary debunker’s problem doesn’t really have anything to do with our knowledge of evolution, or anything else about the formation of our moral inclinations. As noted before, it is not as though we have any reason to suppose that natural selection is somehow *anti-reliable* with respect to moral truth. The debunker’s only plausible claim is that evolutionary forces operate *independently* of moral facts. At worst, evolutionary considerations *fail to vindicate* our moral judgments, they don’t provide any *further reason to doubt* the reliability of these judgments. Street talks of the *distorting* influence of evolutionary pressures. But this way of putting it can be misleading. Usually talk of distortion assumes a background of general reliability. (A signal that fairly reliably transmits information can be distorted). I can do a decent job of assessing job candidates but unconscious biases with respect to race and gender might distort my judgment. Evidence of such distorting influences may lower my estimate of my reliability and thereby lower my confidence in my judgment. But evolutionary forces don’t have a general tendency to *lower* the reliability of moral faculties over generations. It is not as though there is reason to think that our ancestors from microbes to monkeys had a better grip on moral reality which got distorted down the line. Natural selection just prunes branches off the tree of life by mechanisms causally independent of the moral facts. This at any rate is the most plausible story that could be used to motivate skeptical problems. But if this is so then evolution is providing no skeptical problem that we didn’t already have.

Taking the cue from Quine, we might think that at least our moral faculties are worse off in this regard than induction, perception, memory, and other sources of belief whose reliability has more obvious survival value. But on reflection it is not plausible that this could make a significant epistemological difference. Imagine our friend Bill who while otherwise much like us is misled into doubting standard claims about evolution. Given certain misleading evidence he is not justified in supposing that we are the products of many generations of

random mutation and natural selection (and similarly he has no idea whether we were created by a benevolent God, or anything else). Should Bill expect the sun to rise, or think that he has two hands? For that matter, should he even be even a little less confident on these matters than the rest of us? Even as skeptical views go, the following is especially implausible:

Most of us educated folks are fortunate enough to have an adequate basis to judge that our prediction that the sun will rise tomorrow is likely to be right. That's because we've now learnt that we're the products of a process that is likely to favor folks with the ability to latch onto those regularities that our world happens to display. But poor Bill has no such basis to suppose that any of his faculties are reliable as he is ignorant concerning their origin. (Of all the *possible* cognitive systems one could have, how could he have any idea whether he is fortunate enough to have ones that provide him with true beliefs?) So unlike us, he can take no more than a wild guess about what the sun will do tomorrow, whether he has hands, or what he ate for breakfast. (He can't actually help but believe in much the ways that we do (including the belief that many of his beliefs are justified) but in fact his everyday beliefs have a far inferior epistemic status than ours).

Perhaps there is room to sensibly think that evolutionary considerations can add *something* to the epistemic status of our ordinary beliefs. Perhaps there is a kind of virtuous explanatory coherence to a system of beliefs that includes a theory explaining the reliability of our faculties (those faculties on which we relied in developing the theory). But anything much beyond this is just not plausible. Aside from being implausible on the face of it that there is a stark epistemic difference between us and Bill, there is the obvious puzzle about how we ever managed to gain justification to believe evolutionary theory in the first place if prior to that we were in poor Bill's predicament. Unless we accept a more general skepticism we can't plausibly think that the justification of our perceptual beliefs or inductive conclusions *requires* a Quinean vindication from natural selection.¹⁴ But if we give them a free pass, why be so hard on our moral beliefs?

It may be instructive also to compare the case of our more sophisticated non-moral judgments. It is one thing to suppose that natural selection would favor organisms with an ability to track the movements of predators and prey. It is quite another to suppose that there has been a reproductive advantage to being good at linear algebra, or theoretical physics, or evolutionary biology, or epistemological reflections on evolutionary biology. This is not to say that these phenomena are beyond the reach of plausible evolutionary explanation. Perhaps brain structures that allow us these higher intellectual pursuits also provided more mundane benefits relevant to survival. Assuming that we *do* have these advanced cognitive capacities, some such explanation may be plausible. But the plausibility of such an explanation cannot depend on these capacities being *predictable* on the basis of evolutionary considerations alone. Returning to our mystery Planet X, perhaps we should expect that those creatures whose ancestors have survived over

generations there will do a decent job at basic logic, arithmetic, and induction. But we can't possibly predict that they what intellectual heights they will attain dramatically beyond that. Even on Earth a small proportion of creatures have a strong aptitude for scientific or philosophical theorizing. Most people who fancy themselves well equipped to evaluate biological or epistemological hypotheses make a complete botch of it. There is no telling what intellectual nonsense evolution may throw up.

Once again, perhaps the case of moral judgment is even worse off here. The prospects of explaining how (some of us) came to be able to reliably theorize about our evolutionary origins may look better than explaining how we came to be so good at ethics. It is hard even to get a grip on what odds we might give the products of natural selection being reliable at ethics versus reliable at evolutionary biology. But perhaps if we had to take a guess at one or the other, it would be safer to guess that the creatures on Planet X did a decent job at evolutionary biology or epistemology than moral judgment since with respect to the former we can at least glimpse how an evolutionary explanation of our reliability might go. But now if evolutionary considerations are supposed to be our guide in assessing our cognitive reliability, this will only allow us a rather feeble degree of confidence in advanced science over ethics. But that won't do for the aspiring debunker. Any debunking project that appeals to the results of evolutionary biology and applies epistemological reflection to relevance of these to the status of morality had better think that science and epistemology are on relatively solid footing while our moral judgments are built on sand. We are just not going get that stark a difference out of evolutionary considerations. Rightly or wrongly we judge that at our best we are pretty good at science and epistemology in a way that goes well beyond what we might expect on the basis of evolutionary considerations. We have yet to see a reason not to treat our moral judgments the same way.

An interesting and very different way of posing the epistemological problem is suggested by Josh Schechter (ms.). Certain facts are *striking*, they "cry out for explanation". It is a cost of a theory if it leaves striking facts unexplained. The virtue of the theory of natural selection is that it can account the phenomena that would otherwise seem incredible. We find it incredible that the eye might have popped into existence fully formed by a large single-step mutation. The story of small cumulative changes over generations can render this explicable. So there is a kind of tension in any theory committed to the following kind of triad (i) Fact F obtains. (ii) F is striking. (iii) F came about by an improbable accident. When faced with an apparent conjunction like this there is theoretical pressure to reject at least one of (i)-(iii). And hence insofar as (ii) and (iii) are hard to deny we have reason to doubt (i). Now one such striking fact (if it is a fact) is that we are quite reliable at forming moral beliefs. It would be astonishing if just some random process resulted by an extremely improbable fluke in creatures with the remarkable ability to reliably track the mind-independent moral facts. So by the theoretical considerations sketched above, if we can't explain how we might

have come to have reliable moral faculties, then we have reason to doubt that we do have such reliable faculties. But it is indeed hard to see how there even could be an evolutionary explanation of our coming to have reliable moral faculties. Hence we ought to doubt our reliability.¹⁵

Like earlier attempts to raise skeptical worries from evolution, Schechter's argument does not appear to rest heavily on evolution as such. Even if we were agnostic about our origins and so had no idea how to explain how we might have come to be reliable moral thinkers it seems that Schechter's argument should apply. So there is room for the same concern that the problem Schechter presents, if there is one, should lead to a much broader skepticism. But in fact in to some degree Schechter's argument may get more purchase in the light of evolutionary theory. Perhaps if we are agnostic about our origins we should think that there surely *is* some explanation for how we came to have reliable faculties even if we haven't discovered it yet. In the case of moral judgment, the more we learn about our evolutionary origins the less likely it seems that there even could be such an explanation. And so the pressure to deny our reliability that Schechter suggests is increased.

There is much in Schechter's argument that seems right to me. I'm sympathetic to the idea of certain facts being striking and in need of explanation.¹⁶ The reliability of our moral faculties is a good candidate for such a fact. And I think there is a kind of tension in a (i)-(iii) conjunction. The powerful thought lying behind Schechter's argument, it seems to me, is that it appears "too good to be true" that a random process should conveniently provide us with the remarkable ability to make accurate moral assessments.¹⁷ It is astonishing that this should come about. And astonishing events are always some grounds for suspicion. Nevertheless, I think we can grant these insights without drawing any skeptical conclusions regarding our moral judgments.

Suppose we have a robot that can output strings of arithmetical symbols. The program that the robot runs on is chosen by a random generator out of trillions of possible programs. It would be a striking fact indeed if the randomly generated program is one that reliably produces true arithmetical equations! Suppose that appears to be the case. Examining the printout we see a long string of well-formed sentences which according to our calculations are correct. A puzzling outcome indeed and one that we cannot sit comfortably with. But now the tension here seems to push in just one direction: against the assumption that the program was randomly produced. It doesn't seem that we should have any difficulty in establishing for ourselves that these equations are indeed true. We can just do the math as we usually do. We do have reason to doubt that the program was really just produced by accident. There is a nice Bayesian explanation here: A reliable arithmetical calculator is much more to be expected given that it was *not* produced accidentally. There are alternative hypothesis — e.g. someone has rigged the system for kicks — which are not themselves extremely improbable but given which the likelihood of the robot doing this is much greater. It follows that what was antecedently very likely — that the program was generated randomly —

has been strongly disconfirmed. Arguably it is just this feature — challenging our assumptions by confirming rival causal explanations — that makes this outcome striking.¹⁸

There is a different kind of case that does seem to call for skepticism about the robot's reliability. We set up the experiment and you go out for lunch while I record the results. As you return I say "You won't believe what happened. The robot is reliably generating hundreds of true theorems." I'm right. You won't believe what happened. Insofar as you can assure yourself that the program really was generated randomly you will find the story too incredible. And rightly so. This might seem to support Schechter's conclusion that we have reason to doubt the reliability of our faculties if we can't explain how we non-accidentally came to have them. But here it is crucial that your evidence for the truth of the robot's output sequences is my *testimony*. What you naturally suspect is that I'm pulling your leg. There is a nice Bayesian explanation of this also which relates to the striking character of the alleged outcome. Even if I'm generally trustworthy, if I were to lie about the outcome I am likely to report some striking outcome as these stand out as more interesting possibilities. Since a tiny proportion of the possible outcomes are striking, a report of a striking outcome is far more likely given that I'm lying than given that I'm telling the truth. So my report of a striking event gives you evidence that I'm lying. And you can be no more confident of the truth of what I report than you are that I'm not lying. So you have reason to doubt that the robot is producing true arithmetical theorems even though I'm generally trustworthy. Contrast this with a case in which I claim that the robot produced a particular long string of gibberish.

3%—{1{8* 54+ = = c^88—+>383cxnd%@588—v

A report that *this* very sequence of characters was produced is no more to be expected given that I am lying than that I'm telling the truth. So my report gives you no reason to doubt the veracity of my claim. Since you know me to be generally trustworthy you have plenty of reason to believe that that is what the robot produced.¹⁹

The contrast in plausibility here between the striking and mundane outcomes crucially depends on your source for the claim that the robot produces true theorems being the testimony of another. Once we remove this the contrast disappears. Suppose in an attempt to convince you I hand you the printout with rows of true theorems. Prior to examining them you will remain doubtful that they are true. But you should easily be able to rationally convince yourself that *the statements on this sheet* are true just by doing the calculations. Once you establish that they *are* true, you will have reason to suspect of course that this sheet is not the one produced by the robot but one that I put together to fool you. That is a more plausible explanation of why the sheet contains true theorems. You will continue to doubt my story that *the strings produced by the robot running a randomly generated program are true theorems*. You will just conclude that those

are not these. No doubts are raised about your judgments of the truth of any theorem. Note that this result does not depend on the assumption that you are certain that you are a super-reliable arithmetician and hence that it is extremely implausible that you could be mistaken in your judgments. Perhaps you realize that you are prone to errors yourself and so are cautious in your judgments. But the caution called for here is just that which is called for in your evaluation of any arithmetical claims regardless of their source. I see no additional pressure to doubt your judgments due to the striking character of a randomly programmed robot producing true theorems.

There can appear to be a further worry when I imagine that *I* am the randomly programmed theorem producer. I've just taken a drug that has a one in a billion chance of restructuring my brain giving me the ability to do advanced tensor calculus. By an extraordinary fluke I now find myself whipping through problem sets and getting them all right. I really am grasping the concepts and following the steps just as a mathematician does. But it is natural to think that before I have at least checked my answers in the back of the book I ought to be suspicious that I really am pulling off this remarkable feat, and doubt that my answers are right. But there is no such problem for an external evaluator. *You* can do tensor calculus and you know that you can, having learnt it in the orthodox way. You can perfectly well judge whether my answers are correct. (Of course if you do verify that they are right you will have reason think that they are not really my answers and that some trick is being played. Probably I just copied them out of the book). If there is a problem with me trusting my own judgments then it seems to stem not from considerations of strikingness but a worry about a kind of circularity. The worry must stem from the fact my only grounds for supposing that I have these new found mathematical abilities comes from the use of these faculties themselves. This kind of complaint can be raised about my trust in my moral judgments. But I don't see that Schechter's arguments raise a further worry.

I think we can only conclude that evolutionary considerations have little role if any to play in debunking our moral beliefs in particular. Insofar as there are skeptical puzzles in the vicinity they are ones that afflict a far wider range of our convictions.

4.3. Luckiness

So let's now look more closely at the worry concerning contingency or luckiness in general, and in particular in cases like Cohen's grad school choice or alternative religious upbringings. Cohen notes that if he would have believe differently under the influence of a Harvard education then it seems he must take himself to be lucky to be right, if right at all. For his choice of graduate school was not at all directed in a way conducive to philosophical truth. He could just as easily have ended up with the opposite opinion.

I suspect that at the end of the day most of us are content to live with this sort of contingency in most cases. Gideon Rosen (2001) reflects on the fact that had he been raised by Genghis Kahn he would have a very different moral outlook. Had he been raised by empiricists he would find the case for the atomic hypothesis unconvincing. Yet Rosen is not moved in his conviction that torture is wrong or that atoms exist, nor does he think he should be. Interestingly, Rosen suggests that the case of theism might be different. If not compulsory, Rosen suggests that it is at least admirable for a theist to reconsider his beliefs in the light of the contingency of their origin. What could make the difference in these cases? Rosen emphasizes that the theist he has in mind is one “who on reflection takes himself to have no positive grounds for his belief: no arguments, no compelling authority, and most importantly, nothing he would regard as a direct experience of the divine.” (p. 85) Once we make this stipulation it is plausible that what drives Rosen’s case is just the theist’s *lack of grounds* for his belief, or perhaps his taking himself to have no such grounds. It is doubtful that the contingency of his belief has any work to do here.

What is the epistemological problem posed by luck? Certain kinds of luck have been thought to be incompatible with *knowledge*. A natural way to diagnose the Gettier examples (both Gettier’s own and the many variations) is that the subject lacks knowledge because his belief is only correct by *accident*. Smith is justified but mistaken in thinking that Jones will get the job and that he has ten coins in his pocket. He naturally thinks that whoever will get the job has ten coins in his pocket. And he happens to be right, since unbeknown to Smith he himself will get the job and just happens to have ten coins in his pocket. His belief is true but due to the quirky causal set up it is just a *fluke* that it happens to be true. Brian Leiter (2004) suggests that the Gettier result and the subsequent naturalistic approach to epistemology (in particular the idea knowledge depends on the etiology of belief) is crucial to the epistemological relevance of the etiological theories of Marx, Nietzsche and Freud.²⁰

I don’t think that this is the most promising way to think about it. First, the lesson from Gettier is supposed to be that justified, true belief is not sufficient for knowledge. That is, if we are convinced by Gettier’s examples we judge that Smith for example lacks knowledge while lacking nothing in the way of justification. He has perfectly good evidence strongly supporting his conviction, he ought not change his opinion one iota (without obtaining further evidence). A *blocking* debunker, who claims that your being justified is incompatible with the causal facts certainly can’t appeal to the Gettier results. Perhaps blocking *knowledge* is an achievement none the less. But while it is not trivial, this kind of skeptical result can seem rather anemic on its own. The more full-blooded upshot that most debunkers have in mind is at the very least that we *ought* to lose confidence in our beliefs if not lose them altogether. Gettier-inspired blocking of knowledge doesn’t by itself have any such consequence.

The more promising strategy here is that of the *undermining* debunker. I take it this is what Leiter has in mind: “To be sure, beliefs with the wrong etiology

might be *true*; but since they are no longer cases of knowledge, we have no reason to presume that to be the case.” (p. 104) Of course it can't be that my awareness of *someone else's* lack of knowledge is enough by itself to undermine my justification. I'm very often justified in believing things that I know others are ignorant of. But perhaps I can't be justified in believing that *p* if I know that *I* lack knowledge that *p*. This can't be quite right. I can't *know* that I won't win the lottery no matter how many tickets there are — I *might* win; someone has to win and it could just as easily be me as anyone else. But even though I know that I don't know that I won't win, I'm still justified in being about as certain as I am about anything (including many things that I do know). In the case of Gettier examples it may be that when a subject learns of his causal predicament (without also learning that his belief is nevertheless true) he thereby loses his justification. Smith for example might learn that Jones will not get the job. He can now see that his belief that whoever gets the job has ten coins in his pocket is at best true by accident. For it will be true not by virtue of the basis of his conviction but by someone else just happening to both get the job and have ten coins in his or her pocket. And he surely can't rationally maintain his belief in this case. But there is a more straightforward explanation for why Smith loses his justification in this case: he based his belief on the assumption that Jones would get the job, but he is no longer justified in thinking so. Since he has no other evidence suggesting that the man who will get the job has ten coins in his pocket, he is no longer justified. The correct diagnosis of the defeat of his justification has nothing essentially to do with etiology and knowledge. None of this is to say that information about 'bad' belief causation has no potential to undermine justification. But considerations of knowledge and Gettier cases are largely a distraction. We will do better to consider more directly what relevance etiology can have to justification.

The contingency of origins of belief, and the luck of being right are not as such any grounds for doubt, as familiar examples show. If I hadn't studied philosophy I would not believe that Hume was born in 1711. I would, if not disbelieve it, give little credence to that particular year being his birth date. And in fact I just learnt this fact by randomly flipping open one of many books on my shelf and reading where my finger landed. I was lucky indeed to be right on this matter! Of course there is nothing unsettling about this. There is nothing problematic about being lucky in obtaining *evidence* for one's belief. Other cases can seem more murky. Suppose I learn that I was adopted from an orphanage. Had my adoptive parents arrived a day later I would have been raised in a family of *counterinductivists*.²¹ I would have been raised to draw the opposite of standard inductive conclusions. (From the sun having risen every day I would conclude that it will not rise tomorrow). It is a matter of luck that I reason in the regular way. I doubt that any of us would really come to doubt the cogency of our reasoning upon learning this fact about our past. That we are more inclined to doubt in cases like Cohen's or our own political or religious beliefs says more

about how questionable we feel they are quite apart from any facts about our nurturing.

It is tempting to dismiss all worries concerning causal contingency as entirely irrelevant to the justification of certain beliefs. This seems to be Ronald Dworkin's (1996) position with regard to ethics. Very roughly: Dworkin dismisses skeptical worries about moral beliefs that appeal to causation as *non sequiturs*. Our reasons for moral conviction come from within the realm of morality. My justification for say, opposition to the war will come from moral principles. It does not involve a causal-explanatory inference from empirical facts as in science. So any business about causal connections or my moral upbringing or evolutionary history is beside the point.

This feels far too quick. The following story will hopefully illustrate why.

Coin in the Head: You are engaged in some extended moral reasoning. Let it be as sophisticated as you like. Perhaps you are weighing general principles against particular judgments in order reach reflective equilibrium. We cut open your head and poke around in there finding the various neural pathways involved in your current reasoning. (You are too engrossed in thought to notice). What we find is that the wiring all leads back to a box. In the box is a coin that gets tossed around. We reach in with a stick and flip it over and find that all of a sudden you change you position on the ethics of torture. As we turn it back and forth you keep flip-flopping. It turns out that all of your moral judgments go via this coin in the box. And our best scientific theories suggest that the coin is being tossed in an objectively random manner.

Can you really continue to have confidence in your moral judgments in this scenario? At any point you can provide a very subtle defense of your position. Your moral convictions form a well-integrated web. If you were to spell out your grounds for your moral convictions they would not include hypotheses about the structure and workings of your brain. Can you just dismiss the empirical findings as irrelevant to the moral question you are considering? But you know that it is just a coin that is responsible for it all. It seems that you cannot continue to believe that the coin is really falling by chance while continuing to trust your own moral judgments. Suppose you have some important moral decision to make where the right thing to do is not immediately obvious. You could choose to engage in moral reasoning to make a decision. Or you could just flip a coin. But you know that that is all that's going on in your head anyway. It's just a choice between coin in the hand or coin in the head. Why bother giving yourself a headache by thinking the problem through carefully? But of course it is absurd that you would make important moral decisions by the flip of a coin.

Once we are disturbed by the thought of a coin in the head, we can bring it outside the head for what looks like the same effect. Instead of the coin being flipped for each moral question, the coin determines the cognitive architecture of your brain, which determines what moral judgments you make. Is it any

more comforting to have the coin operating outside the head? But now we have reached the case of grad school inculcation. Which grad school you go to helps determine your moral outlook which determines the moral judgments you make. And perhaps you chose your graduate school by a coin flip. This appears no better than the case of the coin in the head. If in the end your opinions all trace back to coin toss, what does it matter where in the causal chain we locate the coin?

However, there is reason to be cautious about moving the location of the coin. A large part of what is disturbing about finding the coin in my head is that it seems I can't really be engaging in moral *reasoning* at all in this case. While it's hard to say what moral reasoning consists in, random coin tossing is not it. I'm not responding to relevant factors. I'm not really appreciating the force of various considerations. All of that is an illusion. Here the *blocking* debunker has his day. The discovery of the coin shows that I never have been justified in any of my moral convictions. Once the coin is located outside my head the matter is not so straightforward. Even if the coin toss determines which moral upbringing I receive which in turn determines the conclusions I draw, the reasoning I engage in as a result of the toss and subsequent nurturing may be perfectly cogent. The blocking debunker doesn't get a foot in here. Other things being equal, sound reasoning results in justified belief. It is entirely possible for chancy events to result in a process of sound reasoning.

I will toss a coin now. If it lands Heads I'm going to come and kidnap you and force you to join a cult that teaches that President Obama is a disguised alien bent on colonizing the Earth from Mars. If it lands Tails I'll leave you alone. Your future political beliefs depend on how this coin lands. You will be lucky indeed if you end up believing something remotely true. Phew! It landed Tails. (If it had landed Heads I couldn't make my philosophical point). Your subsequent political judgments have not been impugned in the slightest by my experiment. No doubt this case fails to capture the distinctive details that make some cases of chancy influences on my beliefs seem problematic. But we should be careful not to be misled by the thought that the relevance of a coin toss to my beliefs is same regardless of what role in the causal chain it plays.

Returning now to our graduate school hero, we said that Cohen could *very easily* have come to the opposite conclusion had he attended a different school. This may add some drama to the story by emphasizing his *luck* (What a close shave! Once false move in enrolment and he would have wound up philosophically mistaken!) But it is of dubious epistemological relevance. Perhaps he couldn't have gone to Harvard. Perhaps he was blacklisted in the United States for his communist sympathies. Perhaps some deep equations in fundamental physics reveal that he couldn't so much as have existed without subsequently going to Oxford. Not a very plausible supposition, but even if we could suppose it, that there was no real *risk* of his having formed different philosophical opinions seems to do nothing to alleviate the apparent problem Cohen faces.

The real problem lies in his apparent inability to tell whether the factors determining where he attends grad school — be it an arbitrary choice, political

forces, fundamental physics — put him in a position arrive at the truth. That is clearly the position he is in *before* he packs his bags for Oxford, at least if he is aware of the way that schooling will affect his opinion. Suppose he has been told in advance: if you go to Oxford you will come out believing p ; otherwise at Harvard you will believe $\sim p$. Regardless of what causal processes were at work he can't have clue at this stage as to whether he will end up with a true belief on the matter. This has nothing to do with nature of the causal processes leading him to one school or another. It is just that there are two paths only one of which leads to the truth, but he has no grounds on which to tell which one it is.

That's his predicament prior to grad school. What about after? Of course perhaps he is exposed to poor arguments which he is deluded into thinking support his conclusion. In this case he won't be justified in supposing that he has arrived at the truth. But this will have nothing to do with his peculiar causal predicament. It will just be an unfortunate case of poorly supported belief regardless of how he came by it. To consider the distinctive worry raised by his predicament we should suppose that he is fortunate enough to be in the best possible outcome. We can imagine that in going to Oxford he is exposed to the full range of considerations available, that these do indeed support the conclusion that p which Cohen draws on this basis. The problem has to be that even in this good case Cohen isn't justified in supposing that he has reached the right conclusion. But why not? It appears that he is much better off than he was prior to grad school. He now has excellent evidence for p . Doesn't this justify him in believing p ? But if so, then since he can tell that p is what he believes then he has every reason to suppose that he has arrived at the truth (even if he could not in tell advance that he would be so fortunate).

There does *appear* to be something fishy about the line above. One stab at what is wrong is to suggest that even if Cohen does *in fact* have good reasons for his belief, he is in no position to tell that this is so. We need to be careful here not to conflate the question of what Cohen is justified in believing and what he is justified in believing that he is justified in believing.²² To get to the conclusion that he shouldn't believe p we will need to appeal to an epistemic level bridging principle along the following lines:

JJ-thesis: S has justification to believe p only if S has justification to believe that he has justification to believe p .

Principles like this should raise alarm bells concerning *infinite regress*.²³ And more generally principles of this form have been much debated.²⁴ But even if we set aside such worries it is not immediately obvious why Cohen can't tell that he is *now* justified in believing p . He couldn't have predicted this in advance. But that was before he was familiar with the arguments supporting p . We are not supposing (or arguing for) a more radical conclusion that we can never be justified judging that philosophical argument is any good. We can suppose that Cohen is generally quite competent at this. Why can't he examine these arguments

and see that they do support p . He could think “I had no idea whether I would end up justified in my opinion about p . But luckily for me I am. For these arguments do support p which is my conclusion.”

Perhaps the following story will help capture our uneasiness with this response.

Lucky Moore: There are two doors into a room which contains either a *hand* or a *banana* (not both). If one enters through Door A there will appear to be a hand there; if one enters through Door B there will appear to be a banana. (At one of these doors stands an invisible demon who will induce a vivid hallucination as you enter). G. E. Moore enters by Door A having no idea yet what is in the room or at which door the demon is lurking. Luckily for him the demon was at Door B and Moore perceives the hand that is there. “Here is a hand!” thinks Moore. “Apparently I escaped the demon since my judgment here is correct.”

It does seem absurd to suppose that Moore is in any position to tell whether there is a hand there. Once he headed to Door A he knew it was going to *seem* to him that there was a hand there (just as it would have seemed to him that there was a banana had he taken the other door). This appearance was as likely as not to be veridical. And surely it still is from his perspective, even though he really is looking at a hand. What Moore needs apparently is some evidence concerning veridicality of his experience which is *independent* of that experience. He could discover that hands are more common than bananas in these odd experiments, or that demons prefer to lurk behind doors on the left than the right. Such information can support his judgment. But a judgment that is derived from that experience itself like ‘Here is a hand’ seems of little help when the veridicality of that experience is what is in doubt.

It is tempting to extent this sort of conclusion to similar cases involving not perception but rational inference.

Lucky Russell: There are two doors into a room which contains the premise E of a logic problem from which either p or $\sim p$ can be derived (but not both). If you enter through Door A then upon calculation you will reach the conclusion p ; if you enter through Door B you will conclude $\sim p$. (At one of these doors stands an invisible demon who will mess with reasoning capacities and lead you to a false conclusion that seems obviously correct to you). Bertrand Russell enters by Door A having no idea yet which conclusion does follow or at which door the demon is lurking. Luckily for him the demon was at Door B and by impeccable deductive steps he reaches the conclusion p . “ p ” thinks Russell. “Apparently I escaped the demon since my judgment here is correct.”²⁵

Is Russell in any position to tell that he reached the correct conclusion when he knows that had he entered the other door the opposite conclusion would *seem* just as right? It is natural to think that what he needs is something *independent* of the reasoning that lead him to p on which to judge that he has reached the

truth. This seems to be crucial to the problem for Cohen. In judging whether his conclusion is right all he has to go on are the reasons that lead him to that conclusion. He notes that the superior architecture of Oxford is hardly evidence that they do better philosophy. This would be independent of his philosophical judgment in the appropriate way, but unfortunately doesn't help his case.²⁶

I think there is reason to be cautious about the parallel between the Lucky Moore and Lucky Russell cases. A natural way to diagnose Moore's case is as follows. All Moore has to go on in judging that there is a hand in the room is that there *appears* to be one. Normally the appearance of a hand would be excellent evidence that there is a hand. But for Moore in this case a hand appearance is no less to be expected given that he is hallucinating than given that he is not. For whether hallucinating or not he may appear to see a hand or a banana, and has no reason to expect one over the other in either case. And that there appears to be a hand can hardly be evidence that he is not merely hallucinating a hand. To carry this diagnosis over to Russell's case would be to say that all he has to go on when he enters the room is that it *seems* to him that the answer is *p*. And yet in this situation the fact that it seems this way to him can provide no evidence that this is the right answer.

But this is a dubious way to think of the epistemology of rational judgment. Normally when I draw some conclusion from premises, whether by application of logic, statistical inference, or on the basis of the explanatory virtues of a theory, or what have you, I don't take my evidence to consist solely in the fact that my conclusion *seems* right to me.²⁷ My evidence consists in the premises from which I reasoned. I'm justified by virtue of having taken rational steps from premises I had reason to accept. By hypothesis Russell is given evidence *E* from which he competently deduces *p*. To suppose that he is not thereby justified in his conclusion we must claim that Russell is required to "bracket" the evidence and line of reasoning which led him to it and assess his likelihood of being right on grounds apart from this.²⁸ But this it seems does not accord evidence its proper role. Note that in ordinary cases of reasoning without the threat of demons it is always possible that I get confused and draw some irrational conclusion that seems perfectly correct to me. In advance of doing a bit of arithmetic I can't be too sure that this won't happen since after all it has once or twice before. But is the confidence I should have in my answer *solely* a function of an estimate I can make of being right independently of the evidence I have for my conclusion? It is a truism that one's judgment should be based on one's *total evidence*. Even if the demonic threat should temper Russell's confidence in some way, has the fact that he has strong evidence for his conclusion *no* role to play in justifying his conviction?²⁹

Admittedly it is by no means easy to feel comfortable with the idea that Russell can rationally believe *p*, or indeed have any degree of confidence in his answer in this case.³⁰ One way to bring out forcefully the worry is to raise the stakes.³¹ Suppose you are Russell and are asked to bet at 2:1 that your answer upon entering the room will be correct. Of course you will decline the bet. It is

as likely as not that you will be wrong, and you stand to lose twice as much as you could win. Now you take your chances through a door and are given the question, “ $34 + 27 = ?$ ” The answer “61” seems blindingly obvious to you no matter how many times you check it over. *Now* would you like to take that bet? Well why not if you really think your answer is right? Perhaps you think, “*Of course* I’m right in this case. I can plainly see that this is the right answer.” But you can recall the last time you entered this room you were sure that $22 + 36 = 54$. It wasn’t until the demon’s spell wore off that you saw how confused you were. Indeed this has been happening about as often as not when you enter this room (just as you would expect) even though you can’t help but feel “*This time I’ve got to be right!*” Perhaps you’re not so sure about the bet now. We can offer you other bets at more modest odds $x : y$. But if you still feel reluctant whenever $x > y$ this suggests that you don’t really think you have any more reason to suppose you are right than you did before entering the room.

If we think that Russell should have more than $2/3$ confidence that his answer is right when his evidence does indeed support his conclusion then apparently he should take the bet. But now we can put him through the experiment repeatedly. The demon randomly chooses a door and Russell is given a different question each time he enters. He had better not bet each time he *thinks* he’s right, as that will be every time. And about half the time the demon will delude him into thinking he’s right. He will lose money very quickly this way. Perhaps the policy that he *should* follow (and which will indeed win him money!) is to bet on all and only those occasions on which his conclusion is justified. But it is clear that he just can’t follow this consistently (or at least try as he may, the odds of him succeeding are very slim). Even if we suppose that on those occasions that the evidence does support his conclusion he can tell that it does, the trouble is with those cases in which it doesn’t. Any *attempt* at taking just those bets that are warranted by his evidence will have him taking every bet and losing half the time. What is he to do but refrain from all betting? But if he shouldn’t bet at all doesn’t this suggest that he shouldn’t think he is right on any occasion in the room, even if the evidence happens to support his conclusion?

But the situation is even trickier. One thing is clear: to avoid losing money, prior to entering the room Russell should try to prevent himself from taking any bets when he is inside. The expected value *now* of his betting when he enters is negative since he is as likely as not to be deluded by the demon. Perhaps he should try to bind himself with the policy of refusing all bets. But this is consistent with his judging that there is an even chance that his future decision to decline the bet will be irrational. He may rationally choose to prevent himself from taking future bets while acknowledging that depending on which door he enters he may *then* rationally look back and think “It made sense at the time to resolve not to take this bet, but now I can see that I should.”

While certainly strange, this sort of epistemic asymmetry is perhaps not intolerable. As others have noted, although I often dream I can tell perfectly well that I’m not doing so now. Dreams just aren’t like *this*. Of course when I am

dreaming I'm somehow unable to recognize the difference.³² Similar asymmetries in rational decision are not out of the question. When you are sober you are in the right frame of mind to tell that you are. But if you think there's a good chance that you will drink too much at the party tonight then it is a good policy to give your keys to someone and insist that they drive you home despite your protests that you are fine. You know if you do drink too much your ability to tell that you have will be impaired. But if you do drink responsibly you may well be able to tell that you are fine and rationally override your previously sensible policy.³³

I suspect that no line we can take on these matters will be entirely satisfying. But it is time to step back and see where this leaves us. In particular, are the difficulties raised here unique to the worries raised about causes of belief in cases like Cohen's, or do they really hinge on points that lead us to a far more general skepticism? The key driving worry in the Moore, Russell, and Cohen cases was something like: Had you been in the other situation, for better or worse, you would have *thought* you were right. You would cite various reasons for your different conviction and it would *seem* to you that these considerations were rationally persuasive. But you can't be right in both cases. But of course that is just our predicament *any time we believe anything*. Whatever I believe, there is always the possibility of my having believed otherwise. And had I done so I would no doubt think that I was right. This point alone can hardly lead to a *limited* skepticism. Similarly the requirement of *independence* is familiar from traditional skeptical arguments.³⁴ It is hard to see how we can assess the reliability of our inductive methods without appealing to their track record of success. It is hard to see what non-perceptual grounds we could have to suppose that our perceptual faculties are reliable. Whenever we consider our most fundamental methods or epistemic standards we can feel that they require some kind of endorsement from the *outside*. But we run out of places to stand. We can't step outside of all reasoning, as it were, to assess whether any of our reasoning is any good. It is very hard to avoid a very general skepticism without having to admit that we can sometimes endorse our reasoning *from within*.³⁵

Philosophy itself is a case in point. There may be a number of reasons to be cynical about the success of philosophy quite apart from Cohen's worries. But I assume anyone taking the time to read a philosophy paper thinks that some of us can do a decent job at philosophizing, whatever that might be. Now the Oracle tells us of a Planet X with creatures in many ways similar to us. They have evolved by similar processes, have similar technological achievements, and so on. And they at least have the capacity to *ask* philosophical questions. How good do we think these folks are at philosophy? Are they about the same as us on average, or much better, or much worse? I take it the answer is that we haven't much of a clue. It is very difficult to assess philosophical aptitude by non-philosophical means. That's why we spend so much time reading writing samples and letters. GREs and GPAs don't indicate much. If we could read some of these aliens' work we might say "This is great." Or "This stinks." But otherwise it's

anyone's guess. But now the Oracle reveals: "The Planet X I was talking about is Earth. The philosophers are you folks." Now we must think "Well in that case I suppose they're not so bad." or whatever it is we think of ourselves. But how do we judge this about ourselves? Apparently not on the basis of considerations wholly external to philosophy. The best we can say is that somehow by being immersed in the practice of philosophy we are in (at least a better) position to assess our achievements.³⁶ If the key idea driving our worries in Cohen's case is his lack of *independent* means of assessing his success, then it is not clear how far it can take us beyond our more general anxiety about philosophy.

4.4. Back to Disagreement Again

Our attempts to locate the source of the epistemological worry arising in Cohen's case have only turned up points that lie behind more general skepticism. I don't suggest that we jump to the conclusion that we should be no more worried about such cases than we are about whether there is an external world or if the sun will rise tomorrow. In the case of perception, commonsense suggests on the one hand before Moore enters the room he is entirely justified in raising his arms and thinking "Here is one hand; here is another" even if he has no non-question begging argument that he is not hallucinating. Commonsense also suggests that once he enters the room he loses all such justification and no amount of insisting that there really is a hand there (true though that might be) can help him. I think it is harder than you might expect to say exactly why Moore should be skeptical in the one case but not the other.³⁷ It only gets harder when we try to say why Russell or Cohen should be skeptical in their cases without this extending to everything they think about. But we still might think that there must *be* some problem over and above general skepticism in Cohen's case, even if it is hard to clearly identify. What I want to suggest now is that whatever else is added, if anything, really derives once again from the phenomenon of disagreement. I'll pursue this line by considering a series of cases.

Exclusive Grad School: Oxford has the only philosophy graduate school in the world. Graduates tend to come out believing p , which is indeed what is strongly supported by the full range of philosophical considerations. Cohen comes to this conviction too on the basis of these arguments.

I take it Cohen has no reason here to doubt that he has reached the right conclusion.

Alternative Physical States: A Super-Physicist reveals to Cohen: "Before you were born the universe was in microstate S , which by the physical laws evolves into a state in which you believe p . There is another possible state S' that the universe could have been in which lawfully evolves into a state in which you believe $\sim p$."

Cohen now knows that he was bound to believe p given the prior state of the universe but if it had been in state S' instead he would have thought otherwise. Of course Cohen is hardly in a position to judge on the basis of the physical details which of these states is more likely to lead him to a true belief. It won't help if the physicist spells them out in massive detail [such and such particle is here with this charge. . .]. Antecedently he could only take a wild guess as to which state will lead to him believing correctly. At best what he can say now is, "Well since p is true [for such and such reasons] then I suppose S must be the state that leads me to the truth." He has no *independent* means of assessing the epistemic benefits of one microstate over the other. But now if determinism is true then that there are two such states just follows from the fact that is *possible* for Cohen to have believed $\sim p$. Facts about physical determinism don't raise skeptical worries on their own. (It's not as though the world's being more *chancy* is grounds for more confidence in the deliveries of our faculties). But then the physicist has apparently added nothing for Cohen to worry about beyond the *possibility* of his having believed something different. There is no distinctive skeptical worry raised here.

Possible Grad School: It occurs to Cohen one day that while there is no other graduate school but Oxford, there *could* have been. Indeed there could have been one with a certain characteristic: students who go there tend to come out believing $\sim p$. It is surely possible for there to be such a school. We can imagine it realized in a number of ways: the students are brainwashed, infatuated with their professors, or the like. (That schools are actually like this is not so plausible. But we are only talking about possibilities here). Now if such a school is possible then it is also possible for Cohen to have gone to such a school in which case he would have believed $\sim p$. Here then is a possible course of history such that had it obtained, Cohen would have ended up believing $\sim p$.

Once again, nothing new is added here beyond the mere possibility of his thinking differently. Cohen has no new information to go on here. He is just reflecting on possibilities. The real worry has to emerge somehow when he discovers that there *actually is* a grad school of this sort. But not just any way of learning this will do. If physics reveals that the universe is infinite then there is almost certainly a graduate school somewhere with the property of producing $\sim p$ -believing graduates. Could this be enough to raise the worry?

As I argued earlier, it is a mistake to think there is any special significance to the fact that in Cohen's actual case he could very easily have gone to Harvard, or that he was very likely to. Consider

Magic Grad School: For kicks, Cohen asked his magician friend to create a philosophy graduate program with just the following specification: if a student such as Cohen were to attend there they would graduate believing $\sim p$. The magician obliges and calls the school 'Marvard'. Cohen even flips a coin to decide between Oxford and Marvard, but ends up at Oxford and a conviction in p .

Cohen now knows that there is an *actual* grad school which he almost attended such that if he had he would have believed $\sim p$. But it still seems that nothing crucial has been added here beyond the possibility of his believing differently. That he could easily have gone to Marvard is irrelevant. It makes no difference whether Cohen used a biased coin or even seriously considered attending Marvard. And if it is possible for there to be a school like Marvard then a magician could make an actual one. But it can't be that easy even for a magician to create skeptical problems. I'm pretty sure it's over 70° outside since it is still August and sunny outside. Normally if I were to see a thermometer reading 67° that would give me significant reason to doubt my initial judgment. Of course it is possible for there to be a thermometer reading 67° , either because that is the temperature, or because it is broken, or because it is a reliable thermometer which happens to be malfunctioning in this rare case. I ask Cohen's magician friend to create a thermometer reading 67° . Zap! There it is. Obviously this has given me no grounds at all to doubt that it's over 70° out. Thermometers manufactured in the usual way are typically quite accurate. There's no reason to expect that magically produced one will be when the only specification I gave is that it read 67° . It could be that it is 67° out and all the magician had to do is give me a regular thermometer. But he could just as easily have given me one that was malfunctioning. His having given me a thermometer at all is no reason to suppose he did it one way or another. And since it most likely *is* over 70° he most likely did the latter.

Suppose Cohen has graduated and forgotten all about Marvard. We are supposing that at Oxford Cohen was exposed to excellent reasons to believe p and he graduates justified in this belief. Now he recalls that the magician had created this other school. It *could* be (as far as Cohen can tell) that his Oxford education has deluded him and all the magician did in creating Marvard was to make a grad school that teaches students to reason well. But the magician could just as easily have created a school where the faculty are nuts and give their students drugs to make them believe $\sim p$. There are any number of ways that the magician could make Marvard to the given the specifications. That Marvard exists is no reason to suppose it is one way rather than another. But since all Cohen's evidence suggests that p is true, he should judge it most likely that Marvard does not train its students to reason well.

Real Grad School: Cohen discovers that there is a regular graduate school namely Harvard which was not created by a magician but developed in the usual way. Their graduate program hasn't started yet. But the Oracle tells him that if students were to go there they would graduate believing $\sim p$.

This is now very close to Cohen's actual case. I'm supposing that no grad students have actually been through the Harvard program in an attempt to separate the epistemological issue here from that of *disagreement*. But it actually seems not to make a difference. Once we suppose that students *would* graduate believing $\sim p$,

actually sending some through and having them graduate as predicted doesn't pose any further worry. So the case presents an epistemological challenge to Cohen's belief no more or less than his actual case does. Unlike the previous cases it can seem that Cohen's belief faces a real challenge here, one that goes beyond the bare possibility of his thinking differently or being mistaken. Why? That students would graduate from Harvard believing $\sim p$ most likely has to do with the views of the faculty and how they teach. Cohen knows that graduate programs are taught by professors that are typically smart, well-informed, and so forth. It is likely that they believe $\sim p$ and can present a strong case for it. And this is a reason to suppose it might be true. How much Cohen should be swayed by this will depend on the issue of independence. We may still argue that Cohen's evidence supporting p is enough to largely outweigh the force of expert opinion. But the case cannot be dismissed as easily as the previous ones could be.

But now we have really just come back to the issue of disagreement. What is driving the distinctive worry here has nothing to do with facts he has discovered about the etiology of his own beliefs. It is just the fact that he has evidence that there are apparently very smart, well informed philosophers who differ in their opinions.

5. Choosing to believe

I want to finish by considering some related issues raised in a fascinating essay by Daniel Garber (2007).³⁸ Pascal (1991) famously advocates choosing to instill theistic belief in oneself not on the basis of evidence but on pragmatic grounds. But following Garber as I understand it Pascal does not think that the faith so obtained is *epistemically* unjustified. The new believer has adequate reasons for his belief not just in the form of some uniquely religious experience but by coming to appreciate the force of evidence that was available to him before he believed. Perhaps he just comes to recognize that "the heavens declare the glory of God; and the firmament showeth his handiwork." (I'm not sure if the Pascal thinks the unbeliever is *irrational* in failing to appreciate the force of this evidence. Perhaps it is more like a cognitive ability that the believer has, like a mathematician who can grasp things and see connections that are opaque to others). Since the unbeliever cannot appreciate these grounds from the outside, he can only come to faith choosing to be open to belief by immersing himself in religious tradition and practice.

Garber writes as an unbeliever who resonates with Pascal's call but is unwilling to take the next step and engage in what would seem to be self-deception. If I understand him correctly, it is not that he is sure that in coming to faith he would be deluded. I take it that Garber sees it as an open possibility that Pascal is right, and that if he took the step of faith he might come not only to believe, but to know that there is a God. In any event, I think it is interesting to consider this kind of middle case of uncertainty. If one is certain that not- p ,

then the prospect of self-inducing the belief that p will seem straightforwardly to be a matter of self-deception. We recoil from it for some of the reasons that we recoil from entering Nozick's (1974) Experience Machine in which we experience endless pleasure while being radically mistaken about our predicament. We value being in touch with reality even if it is not how we would prefer reality to be. In the case of uncertainty that interests me the situation is less straightforward. From the doubter's perspective coming to believe p might involve being led into error, or it might be to make a valuable discovery.

The situation is peculiar, and might seem to arise only in the case of religion. But I have two cases of my own experiences with interesting parallels. For a brief time as a young child I was troubled by what I now judge to be irrational fears. I knew that if I could just reject all such thoughts I would be more content. But I also resisted abandoning such thoughts for the kind of reasons mentioned above. It seemed more important for me to be aware of these possible monsters under the bed (even if I was powerless to stop them) than to be happier but possibly mistaken about my real predicament. But I did end up choosing to banish such thoughts not because I became convinced that they were false, but because it would make my life easier to believe that they were. I soon lost these fears and became sure that they were nothing but irrational fantasies.

There are two questions we can raise about this. First, was there anything wrong with my taking this practical step to influence my beliefs? Second, can I rationally maintain my conviction when I realize that it is the result of a choice I made where truth was not my goal? Whatever we say about the first question it is clear to me now that the outcome was a good one both practically and epistemically. I have no doubt that these early fears were false and irrational. It would be mad now for me to revert to my prior state. I couldn't see that so clearly back then which is why it was only a practical decision that could bring me into a new and better way of seeing things. I don't know that I have any way of assessing whether this move was a change for the better which is independent of the way I see things as a result. But this does little to shake my conviction that I am now right.

Here is the second case. I don't understand much about art, especially contemporary visual art. I've had a number of friends that seem to be enraptured by it. But to me it mostly looks like a bunch of crap. Still, I'm intrigued by it. I have often wondered if there is something there that I'm missing out on, that some people are attuned to certain artistic properties that I'm blind to. Perhaps visual art is for some people like music is for me. In that case I'm envious and would like to be in touch with this aspect of reality that I'm missing. On the other hand, I have the suspicion that much of it is bogus. Perhaps artsy folks are fooled into thinking that they are grasping objectively important features of artworks. Such judgments might be no more than an illusion, a product of the social dynamics of the artistic community.

Despite my uncertainty I have tried to take on the artistic viewpoint in much the way that Pascal recommends attending Mass. I've gone to art galleries and

stared at canvases covered with buckets of splattered paint, hoping to turn into one of those people who finds them profound. In advance I could not be sure what the outcome would be. I might discover “Yes! That’s it. That’s what I was missing!” or I might just be duped into thinking that I have. But I took the risk. As it happens I got neither result. I don’t know if that reveals much about the art. I just got bored and gave up. But it is not obvious to me that there was anything wrong with trying.

Why does it seem problematic to engage in this kind of attempted manipulation of belief? It seems very much to depend on the balance of my estimates of the outcomes. I’m much less comfortable bringing myself to believe that which I’m confident is false than I am at weighing the possibility of enlightenment against the risk of delusion. We take risks of falling into error every day. Go to a baseball game and you are bound to misperceive something going on, make a mistake adding the score, or something. If I were to take on learning quantum field theory (or philosophy for that matter) there is a good chance that would get confused along the way and form some not only false but irrational beliefs. It would be silly to go around with doxastic blinkers on out of fear of catching some false beliefs. It might be thought that the important difference between belief manipulation and open intellectual inquiry is that in the latter a few inevitable falsehoods are outweighed by a much greater number of truths gained. But if it was a balance like this that mattered I could just make a policy of memorizing a page of the telephone directory each day in order to believe lots of truths and make up for that one possibly false belief that I’ve deliberately acquired. I suppose the subject matter of the belief might be relevant. Major questions like the existence of God are fundamental to the way we see the world. We are more reluctant to be wrong on such matters. For unlike baseball scores being mistaken on big questions entails being more profoundly out of touch with reality. But as William James (1909) pointed out, these are also matters that we would like to get *right* (remaining agnostic on the big questions also involves missing out on much of reality). And getting them right involves the risk of getting them wrong.

Garber’s main concern is with the epistemic result of belief manipulation. If he were to go to synagogue and renew his childhood faith, could he legitimately trust his new conviction while knowing that he deliberately chose it?³⁹ The fact that there is a choice involved rather than just an accident of birth surely has no relevance to this question. Whether or not it was responsible of him to engage in this belief manipulation has no bearing on what it is now rational for him to believe. Indeed even utterly irresponsible belief manipulation can result in enlightenment. Having been raised as a Counter-inductivist I had always expected the sun not to rise.⁴⁰ Counter-induction generally led me astray, but of course I just took that as further evidence that it would work the next time. Noticing that Inductivists were a miserable lot I figured their fortunes were probably about to change. So I decided it was time to jump ship and delude myself into the Inductivist mindset. I began a Pascalian regimen of attending Inductivist church

and singing Inductivist hymns. Eventually I came to see that of course the sun is most likely going to rise tomorrow; it has risen every day in the past! I see clearly *now* that my old way of thinking was confused. Of course I couldn't appreciate that back then. And even now I have no way — *independently of Inductivist thinking* — of verifying that despite my irresponsible choice my conversion led me to sanity and truth.

If there is no God then of course Garber will be mistaken and presumably his belief will not be fully rational. If Pascal is right, then he may come to grasp reasons to believe that he had been blind to for years. Perhaps it will be hard not to think at times “But what if I’ve just been deluded into thinking that I’m latching onto the truth?” And he may feel that his belief is in need of endorsement from outside the perspective of his faith. But we have failed to identify a clear reason that such thoughts should concern him beyond the fact that others don't share his view and the general skeptical worry that he *could* after all, be mistaken.

Notes

1. I've learnt a lot from discussion with students and friends who have wrestled with similar questions, in particular Adam Elga, Katia Vavova, and Miriam Schoenfeld. I've presented versions of this material at the Princeton Workshop on Epistemology and Religion, the Pacific APA in Vancouver, the Philosophy of Religion Conference at Sydney University, the St Thomas Philosophy of Religion Summer School, and graduate seminars at MIT and Brown. I've benefited from the feedback I received on these occasions especially from my APA commentators Evan Fales and Tom Crisp. Thanks also to Daniel Berntson, David Christensen, Justin Clarke-Doane, Ronald Dworkin, Dan Greco, Tom Kelly, Tom Nagel, Josh Schechter, Paulina Sliwa, and Sharon Street, for discussion of these and related questions.
2. Street does not recommend moral skepticism but rather a form of constructivism which she argues avoids the epistemological problem. Joyce (2006) presents similar arguments for a more skeptical conclusion.
3. While he doesn't explicitly endorse anything like this, Cohen notes repeatedly that his having gone to Oxford has no relevance to the truth of his view on the analytic/synthetic distinction, giving the impression that this might be what the worry hinges on.
4. The topic has received surprisingly little sustained attention from analytic epistemologists. For work that addresses these issues to one degree or another see Cohen (2000), Dworkin (1996), Elga (ms.), Garber (2007, 2009), Leiter (2004), Plantinga (2000), Rosen (2001), Schechter (ms.), Sher (2001), and Vavova (ms.). On Darwinian doubts about morality see Joyce (2006) and Street (2006). The general idea of etiological theories casting suspicion on common beliefs seems to have played a larger role in continental philosophy. Leiter (2004) has a very interesting discussion of the prominent examples of Marx, Nietzsche, and Freud.

5. And you don't have to be a Moorean to worry that extreme forms of skepticism are self-undermining. See Rinard (ms.) for a response to skepticism along these lines.
6. Some of the seminal papers are van Inwagen 1996, Kelly 2005, Feldman 2006, Christensen 2007, and Elga 2007. Christensen 2009 is a recent survey of the field.
7. Unless he takes it to be what I've called an epistemically permissive case (White 2005b). On this account the evidence does not determine a unique rational doxastic attitude; one can rationally believe p or rationally believe $\sim p$ (but not both) on the basis of the same evidence. But if we really think there are such cases then even meeting an actual disagreeing peer seems to pose no challenge to one's belief.
8. This sort of problem was raised by Jonathan Vogel (2007) against counterfactual tracking accounts of *knowledge* such as Nozick's (1981).
9. It came to my attention recently that some of the points I've made in the last three paragraphs are similar to one's made by Alvin Plantinga (2000). A possible explanation is that my beliefs here can be causally traced back to Plantinga's, either because I once read it or by some indirect route. But this is not the kind of etiology of belief that raises any doubts about their correctness.
10. I'm summarizing (fairly I hope!) some of Tom Crisp's commentary on a version of this paper at the Pacific APA.
11. Some people seem to think that theism is somehow *incompatible* with standard evolutionary theory. I think this can only rest on confusion. See Sober (forthcoming), van Inwagen (2003).
12. Dennett (2006) surveys a range of evolutionary hypotheses concerning religious belief.
13. Street needn't deny that this *is* easy to recognize, she will just give a constructivist account of what this fact consists in and how we can thereby recognize it. Joyce is committed to the more skeptical conclusion.
14. Of course Quine didn't think so either. His explanation was intended just as a potential part of a scientific understanding of ourselves, not a as a necessary answer to a real skeptical threat.
15. Schechter does not focus on the case of moral conviction. But I think this captures the way his argument should be adapted for our purposes.
16. I appeal to this idea in White 2005a.
17. Street also discusses the moral realist's pressure to explain how our ability to grasp evaluative truth could have evolved except by pure chance. But her explicit arguments are quite different from Schechter's. However Street does repeatedly note the extraordinary *coincidence* involved in evolution leading us by pure chance to the moral facts, and how remarkably *convenient* this is for the moral realist. This suggests that perhaps Schechter's arguments lie behind the appeal of some of Street's discussion.
18. This is the analysis that Horwich (1982) gives of what it is for an event to be striking, or as he puts it 'surprising'.
19. This of course is related the problem from Hume (1999) about why we ought to be more dubious about reports of miracles than we are of other highly improbable events like the outcome of a large lottery. That natural selection should provide us with reliable moral judgments can seem like a *miracle*, and thereby worthy of the same skepticism.

20. Leiter doesn't explicitly speak of 'luck' or 'accident' although his discussion of Gettier examples naturally suggests this. At any rate, my dissatisfaction with Leiter's proposal doesn't hinge on whether we think of the relevant feature of the bad belief causes as being their luckiness at getting us to the truth.
21. Adam Elga and Dan Greco have suggested this kind of case in conversation.
22. Alston (1980) complains that epistemologists do this all too often.
23. The falsity of principles like the JJ-thesis is often taken to be the lesson of Lewis Carroll's (1895) story of Achilles and the Tortoise
24. Weatherson (ms.) disputes even the weaker thesis: S is has justification to believe p only if S lacks justification to believe that he lacks justification to believe p . Christensen (2010b) who is sympathetic to higher-order constraints on justification explores some of the difficulties in giving a systematic treatment of such constraints.
25. This case is similar to, and inspired by Christensen's (2010a) examples involving reason-distorting drugs.
26. Katia Vavova (ms.) develops and defends a condition of independence in diagnosing Cohen's and other cases. A similar idea is appealed to in defenses of conciliatory views on disagreement by Elga (2007), and Christensen (forthcoming).
27. Even in the case of perception there is plenty of room to argue that Moore's most basic visual evidence does not consist in facts about how things visually appear to him. See e.g., Williamson 2000. I explore some of the surprising consequences of this position in White (forthcoming).
28. Christensen (2010a) defends the view that it is sometimes appropriate to bracket part of one's evidence in this way.
29. The line suggested here is developed with respect to disagreement by Kelly (forthcoming). I explore the matter further in White 2009.
30. David Christensen has urged in conversation that one needn't be in the grip of any dubious epistemology of "intellectual seemings" as one's only evidence to be persuaded of some independence constraint. The cases may speak for themselves.
31. Another way, developed by Vavova (ms.) is to argue that in rejecting a condition of independence we must allow for a surprising (absurd?) kind of *bootstrapping* over repeated cases.
32. Williams (1978) makes this kind of point.
33. Tom Kelly reminded me that those of us who are Two-Boxers in the Newcomb Problem face some similar puzzles. We think we are rational in taking both boxes even while the One-Boxers consistently make more money than us. Before the predictor makes his decision we may try to turn ourselves into (irrational) One-Boxers by resolving to take one box. But we still think that once the boxes are filled it would be rational to turn around and take the extra box.
34. Elga (ms.) argues that unless Cohen has reason to doubt that he is meeting his own fundamental epistemic standards, the problem raised by his case is just the general skeptical worry that we have no independent support for the correctness of those standards.
35. Paul Boghossian (2006) suggests worries along these lines with respect to our fundamental epistemic rules as a possible motivation for epistemological relativism.

36. Fair minded folks outside our discipline sometimes wonder what it all amounts to. Do we really blame them?
37. Obviously it has something to do with the actual existence of a demon nearby rendering his deception more likely. But this doesn't begin to solve the puzzles in this area.
38. Garber continues the discussion in his 2009.
39. Of course we could run the question the other way just as well. Inspired by Nietzsche's (1954) joyous description of a life without God, I could try to inculcate an atheistic conviction.
40. Unlike the previous cases this one is made up.

References

- Alston, William. 1980. 'Level Confusions in Epistemology' *Midwest Studies in Philosophy* 5: 135–150.
- Boghossian, Paul. 2006. *Fear of Knowledge: Against Relativism and Constructivism* Oxford: Oxford University Press.
- Carroll, Lewis. 1895. 'What the Tortoise Said to Achilles' *Mind* 4: 278–280.
- Christensen, David. 2007. 'Epistemology of Disagreement: The Good News' *Philosophical Review* 116: 187–217.
- Christensen, David. 2009. 'Disagreement as Evidence: The Epistemology of Controversy' *Philosophy Compass* 4: 756–67.
- Christensen, David. 2010a. 'Higher-Order Evidence' *Philosophy and Phenomenological Research* 81: 185–215.
- Christensen, David. 2010b 'Rational Reflection' *Philosophical Perspectives*.
- Christensen, David. Forthcoming. 'Disagreement, Question-Begging, and Epistemic Self-Criticism. *Philosophers' Imprint*.
- Cohen, G. A. 2000. *If You're an Egalitarian, How Come You're So Rich?* Harvard University Press.
- Darwin, Charles. 1998. *Origin of Species* Oxford: Oxford University Press.
- Dennett, Daniel. 2006. *Breaking the Spell: Religion as a Natural Phenomenon* Viking Press.
- Dworkin, Ronald. 1996. 'Objectivity and Truth: You'd Better Believe it' *Philosophy and Public Affairs* 25: 87–139.
- Elga, Adam. 2007. 'Reflection and Disagreement' *Noûs* 41: 478–502.
- Elga, Adam. Manuscript. 'Lucky to be Rational' presented at Bellingham Summer Philosophy Conference, Bellingham, WA August 6, 2008.
- Feldman, Richard. 2006. 'Epistemological Puzzles about Disagreement' in *Epistemology Futures* S. Hetherington ed. Oxford: Oxford University Press, pp. 216–236.
- Garber, Daniel. 2007. 'Religio Philosophi: Some thoughts on God, Reason, and Faith' in L. Antony ed. *Philosophers Without Gods: Meditations on Atheism and the Secular Life* Oxford: Oxford University Press, pp. 32–40.
- Garber, Daniel. 2009. *What Happens After Pascal's Wager?: Living Faith and Rational Belief. The 2009 Aquinas Lecture*. Marquette University Press.
- Horwich, Paul. 1982. *Probability and Evidence* Cambridge: Cambridge University Press.
- James, William. 1909. 'The Will to Believe' in *The Will to Believe and Other Essays in Popular Philosophy* New York: Longmans Green & Co. pp 1–25.
- Joyce, Richard. 2006. *The Evolution of Morality* Cambridge: MIT Press.
- Hume, David. 1999. *An Enquiry Concerning Human Understanding* edited by Tom L. Beauchamp, Oxford: Oxford University Press.

- Kelly, Thomas. 2005. 'The Epistemic Significance of Disagreement' *Oxford Studies in Epistemology* Volume 1, J. Hawthorne and T. Gendler eds. Oxford: Oxford University Press, pp. 167–196.
- Kelly, Thomas. Forthcoming. 'Peer Disagreement and Higher-Order Evidence' in *Disagreement* R. Feldman and T. Warfield eds. Oxford: Oxford University Press.
- Leiter, Brian. 2004. 'The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud' in *The Future for Philosophy* B. Leiter ed. Oxford: Oxford University Press.
- Nietzsche, Friedrich. 1954. *Thus Spoke Zarathustra* in *The Portable Nietzsche* translated by Walter Kaufmann, New York: Viking Press, pp. 115–343.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia* New York: Basic Books.
- Nozick, Robert. 1981. *Philosophical Explanations* Cambridge: Harvard University Press.
- Paley, William. 2006. *Natural Theology* Oxford: Oxford University Press.
- Pascal, Blaise. 1991. *Pensées* P. Sellier ed. Paris: Bords.
- Plantinga, Alvin. 2000. *Warranted Christian Belief* Oxford: Oxford University Press.
- Quine, W. V. O. 1969. 'Natural Kinds' in *Essays in Honor of Carl G. Hempel*. N. Rescher ed. D. Reidel.
- Rinard, Susanna. Manuscript. 'Reasoning One's Way Out of Skepticism'.
- Rosen, Gideon. 2001. 'Nominalism, Naturalism, Epistemic Relativism' *Philosophical Perspectives* 15: 69–91.
- Schechter, Joshua. Manuscript. 'Luck, Rationality, and Explanation: A Reply to Elga's "Lucky to be Rational"'
- Sher, George. 2001. 'But I Could Be Wrong'. *Social Philosophy and Policy* 18: 64–78.
- Sober, Elliot. Forthcoming. 'Evolution without Naturalism' in *Oxford Studies in Philosophy of Religion* Volume 3, J. Kvanvig ed.
- Street, Sharon. 2006. 'A Darwinian Dilemma for Realist Theories of Value' *Philosophical Studies* 127: 109–166.
- van Inwagen, Peter. 1996. 'It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence' in *Faith, Freedom, and Rationality* J. Jordon and D. Howard-Snyder eds., Rowman and Littlefield, pp. 137–154.
- van Inwagen, Peter. 2003. 'The Compatibility of Darwinism and Design' in *God and Design: The Teleological Argument and Modern Science* Neil A. Manson ed. London: Routledge, pp. 347–362.
- Vavova, Ekaterina. Manuscript. 'What to believe when you believe that if things had been different, you wouldn't have believed what you now believe'
- Vogel, Jonathan. 2007. 'Subjunctivitis' *Philosophical Studies* 134: 73–88.
- Weatherson, Brian. Manuscript. 'Does Judgment Screen Evidence?'
- White, Roger. 2005a. 'Explanation as a Guide to Induction' *Philosophers' Imprint* 5.
- White, Roger. 2005b. 'Epistemic Permissiveness' *Philosophical Perspectives* 19: 445–459.
- White, Roger. 2009. 'On Treating Oneself and Others as Thermometers' *Episteme* 6: 233–250.
- White, Roger. Forthcoming. 'Defeasibility and Skepticism' in *Contemporary Perspectives on Perceptual Justification and Skepticism* Oxford: Oxford University Press.
- Williams, Bernard. 1978. *Descartes: The Project of Pure Inquiry* London: Pelican.
- Williamson, Timothy. 2000. *Knowledge and its Limits* Oxford: Oxford University Press.