ON NATURALISTIC METAPHYSICS[†]

Thomas M. Crisp

Blackwell Companion to Naturalism, ed. Kelly James Clark
(Wiley Blackwell, 2016)

Abstract: I raise an epistemic objection to naturalistic metaphysics, the attempt to understand the nature and structure of reality in terms of natural entities, forces, and processes, arguing that we should not expect evolution to have crafted cognitive faculties reliable with respect to recondite metaphysical speculation and that this gives practitioners of naturalistic metaphysics reason to doubt the deliverances of their work. I conclude by considering some main objections to this kind of skeptical argument.

Naturalism—on one popular understanding of the term—is the view that there is no good reason to believe in God or anything relevantly like God, and no need to appeal to such an entity in explaining or understanding the world. No part of the world, says the naturalist, whether it be the mental, the physical or the moral parts, requires postulating non-natural entities, forces, or processes to explain and understand its workings.

Naturalism is highly fashionable in Western academia. Its journals brim with attempts to give naturalistic accounts of the mental, the moral, the physical, and all else besides. In some quarters, especially the sciences, to suggest that there are or might be non-naturalistic explanations of this or that phenomenon is to court professional disaster, and to make oneself the target of vituperative accusations of intellectual sloppiness and irresponsibility. Naturalism, in many sectors of the Western academy, has become orthodoxy.

But there is a deep and little discussed problem with this research program, a problem that threatens the coherence of the whole project. After some further explanation what naturalism is, I turn to an explanation of the problem.

What is naturalism?

Naturalism, I say, is the view that there is no good reason to believe in God (for present purposes, the omniscient, omnipotent, wholly good creator of all), or anything relevantly like God, and no need to appeal to such an entity

in explaining or understanding the world. But what would it be for a being to be "relevantly like" God? A being is "relevantly like God", let us say, if it is a mind (something with beliefs, desires, aims, and intentions) whose activities explain the origin, structure, and ongoing existence of the cosmos.

I said too that, according to naturalism, no part of the world, whether it be the mental, the physical, or the moral parts, requires postulating non-natural entities, forces, or processes to explain and understand its workings. And what are "non-natural" entities, forces, or processes? Entities, forces, and processes are non-natural, let us say, if they cannot be completely described without appeal to the existence of a mind whose activities explain the origin, structure, and ongoing existence of the cosmos.[1]

On naturalistic metaphysics

I suggested above that there is a fundamental problem with the project of explaining the mental, the moral, and the physical in naturalistic terms. It'll help in explaining that problem to say more about what that project comes to.

Appearances can be deceiving.[2] Someone might appear to be a friend but really be a foe; something might appear to be near but really be far; the earth beneath our feet appears to be stationary but is really rotating; and so forth. Sometimes it's possible to get behind misleading appearances and see things as really are. So: when we discovered that the earth is really rotating, we got behind the appearance that the earth is stationary and saw into how things really are.

Metaphysics is the attempt to do this: it's the attempt to get behind the appearances to the reality underlying them. More carefully: it's the attempt to get behind the appearances and understand the nature and structure of the reality underlying them.

We can think of metaphysics as coming in more or less abstract varieties. To the extent that the constituents of the structures postulated by a metaphysical theory differ from the objects of everyday experience, to the extent that the relationships of dependence between the postulated structures and the world of appearances are complicated (requiring complex mathematics, say, and/or long chains of complicated reasoning to describe), to the extent that knowledge of these structures and the dependence relationships connecting them to the world of appearance is far removed from the concerns of everyday life, and to the

extent that questions about these structures and dependence relationships have been the subject of perennial philosophical dispute, thereby the theory in question is more abstract.

A relatively non-abstract form of metaphysics is the sort of inference to the best explanation hunter-gatherers use when tracking wounded animals, postulating hypotheses about the animal's injuries and location from appearances on the trail. Such attempts to get behind the appearances and understand the nature and structure of the reality underlying them are, as we're using terms, a form of metaphysics. Since knowledge of these underlying realities and their relationships to the appearances is highly connected to concerns of everyday life, and the postulated realities and their connections to the appearances are of a familiar, everyday sort, this represents a relatively non-abstract sort of metaphysics.

Not so with, for example, string theory: Here one postulates strange particles, multidimensional spacetimes, and so forth, knowledge of which is far removed from the concerns of everyday life, and whose relations of dependence to the world of everyday appearances are describable only with the help of extraordinarily subtle and complex mathematics. Here we have a highly abstract form of metaphysics.

Such, then, is metaphysics, and some indication of what it is for metaphysics to be more or less abstract. Naturalistic metaphysics, let us say, is metaphysics such that the structures that it postulates to explain the appearances are populated entirely by "natural" entities, forces, and processes: entities, forces, and processes completely describable without appeal to a mind or minds whose activities explain the origin, structure, and ongoing existence of the cosmos.

And such is the project of explaining the mental, the moral, and the physical in naturalistic terms: naturalistic metaphysics. It's the attempt to get behind the mental, moral, and physical appearances, to understand the nature and structure of the realities underlying them, all the while postulating only natural entities, forces, and processes. And as it occurs in the pages of philosophical and scientific journals and books, it's mostly highly abstract metaphysics. The Flanagan, Sarkissian, Wong essay in this volume is an elegant example. It sketches what it would look like to give a purely naturalistic account of moral phenomena, pointing out along the way that there are no irreducible, non-natural moral properties, that there is no such thing as metaphysical

freedom or agent causation, that quantum-mechanical processes in the microtubules of certain neuronal segments would do nothing to secure the existence of agent causation, that there is no such thing as a faculty of the will, and that all moral truths are grounded in facts about human needs, desires, and purposes: these are, one and all, claims about realities far removed from the concerns of everyday life, all subject to massive and perennial philosophical dispute. Some highly abstract, naturalistic metaphysics on display here.

An objection to naturalistic metaphysics

Now for my objection to abstract naturalistic metaphysics (for short, *naturalistic metaphysics*). It rests on three theses. First, there is this

> *Evolutionary Thesis*: If naturalism is true, if there is no need to appeal to God or anything like God in explaining or understanding the world, then it is highly likely that we humans and our cognitive faculties are the product of unguided evolutionary processes of the sorts described by contemporary evolutionary theory,

where by 'unguided', I mean that the processes which gave rise to us were not orchestrated or superintended by an intelligent agent—that they were "blind", to use Richard Dawkins' term (1986). If neither God nor any being like God figures into the explanation how the world works, then neither figures into the explanation how we humans and our cognitive faculties arose, and if that's so, it's extremely difficult to see how we could have got here if not by unguided processes of the sort described by contemporary evolutionary theory.

Next, there is this

> *Thesis of Unreliability*: The probability that we humans command much by way of reliable insight into abstract metaphysical matters, given naturalism and that we and our faculties are the product of evolutionary processes of the sort described by contemporary evolutionary theory, is *inscrutable* (that is, we have no way of knowing its value).

Why think this thesis true? Well, is there some reason, given naturalism and that we and our cognitive faculties arose via

unguided evolutionary processes of the sort described by contemporary evolutionary theory, to expect those processes to have endowed us with much by way of reliable insight into abstract metaphysical matters—matters far removed from the exigencies of normal human life? Surely the answer here is 'no'. According to the usual evolutionary story, human cognitive faculties—of the same basic sort we possess today—appeared during the Pleistocene era, during the period lasting from about 2.5 million years ago to about 12,000 years ago. By the end of that period, our ancestors possessed brains of roughly the same basic architecture and cognitive capabilities as our brains. The main explanation why they evolved those faculties, on the usual story, is that having such faculties was adaptive in Pleistocene environments: useful for feeding, flying, fighting, and reproducing on the plains of Pleistocene Africa. But why would cognitive faculties selected for success at those tasks in those environments have required reliability with respect to abstract metaphysics, reliability with respect to questions about the deep structure of realities underlying the appearances, knowledge of which would've been wholly irrelevant to life on the plains of Pleistocene Africa? From a fitness point of view, such cognitive capability seems wholly unnecessary. But if so, the probability that it should have evolved seems low.

There is, however, the possibility that reliability on abstract metaphysical matters far removed from the everyday concerns of life is a "spandrel"—a non-adaptive byproduct of some adaptively selected trait, just as, for example, reliability on abstract mathematical matters is plausibly thought of as a non-adaptive byproduct of the adaptive ability to do simple arithmetic. That could be, but as I read the cognitive science literature, no one has been able to find good reason for thinking so. More, as I'll try to argue below, the prospects for producing a decent argument for thinking so are dim. The right thing to say, I think, given the current state of the evidence and the considerations I raise below, is that the probability we should have got cognitive faculties capable of reliability on abstract metaphysical matters, given naturalism and the usual evolutionary story, is inscrutable (we have no way of knowing what it is).

Let me pause briefly to note that I am not here recapitulating Alvin Plantinga's famed "evolutionary argument against naturalism",[3] in which Plantinga argues that the probability of our overall cognitive reliability (the probability that our faculties should be generally truth-conducive, with respect even to everyday matters), given

naturalism and evolution, is low or inscrutable, and then goes on to argue that this gives the believer in naturalistic evolution reason to doubt all of the deliverances of her cognitive faculties, including belief in naturalistic evolution. My argument is inspired by Plantinga's argument and similar in many ways,[4] but it differs in this key respect: I am not claiming that, given naturalistic evolution, the probability that our cognitive faculties are generally reliable is low or inscrutable. I am arguing for something much weaker, and, I think, much more easily defended. Objectors to Plantinga's argument have claimed that, contrary to Plantinga, given naturalistic evolution, we should expect our faculties to be generally reliable.[5] For present purposes, I needn't dispute that. I claim only that, given naturalistic evolution, the probability that we'd have got faculties reliable with respect to abstruse metaphysical matters is inscrutable. And that might well be, even if, given naturalism and evolution, we should expect some sort of general cognitive reliability.

Finally, there is this

> *Principle of Reason*: If for some source of information S, you have good reason to doubt the reliability of S, and have no good reason to discount this reason for doubt, then the rational attitude toward matters about which S is your only source of information is doubt,

where the key terms here may be understood as follows. You have *good reason to doubt* the reliability of a source of information S if there is some condition C_1 such that you have good reason to think that S is in condition C_1 and good reason to be agnostic about the probability that S is reliable given that it is in C_1. You have *good reason to discount this reason for doubt* if there is some further condition C_2 such that you have good reason to think that S is in both C_1 and C_2 and good reason to think that S is probably reliable given that it is in both conditions C_1 and C_2. Finally, the rational attitude to take toward some matter is *doubt* if the rational attitude, given your evidence on the matter, is either agnosticism or disbelief.[6]

Why accept the Principle of Reason? I argue by cases. Suppose you have good reasons for agnosticism about the accuracy of a certain pool thermometer (say, you stepped on it and the gauge now rubs against the metal backing), no reason for thinking that, in this case, those reasons don't apply, and no further information about the water temperature in the pool. Then, pretty clearly, you shouldn't

believe what the thermometer says about the water temperature; the rational response here is agnosticism. Or suppose you have good reason for agnosticism about Uncle Fred's reliability on matters of family history (he is suffering from mild dementia), no special reason for thinking that, in this case, he can be trusted, and no further information about the stories he is now relating. The rational response here, one thinks, is doubt. Reflection on these and like cases, I suggest, lends strong support to our Principle.

Our two Theses and Principle in hand, I can now state my objection to naturalistic metaphysics. Suppose you are a practitioner of abstract naturalistic metaphysics, developing this or that purely naturalistic account of the mental, the moral, or the physical, postulating structures and objects quite different than those encountered in everyday experience, connected in highly complex ways to the appearances, knowledge of which is far removed from the everyday concerns of life and subject to perennial dispute. Then, I suggest, you have good reasons for accepting both of our above Theses—the Evolutionary Thesis and Thesis of Unreliability, for reasons I laid out above. But if so, you have good reason to doubt the reliability of those parts of your cognitive endowment responsible for abstract metaphysical beliefs (henceforth, your *metaphysical faculties*). This is because you have good reason to think that your metaphysical faculties are in the condition one's faculties are in if they were produced by unguided processes of the sort described by contemporary evolutionary theory, and good reason for agnosticism about the proposition that one's metaphysical faculties are probably reliable given that they are in that condition. Assuming you have no good reason to discount those reasons for doubt—no good reason for thinking your metaphysical faculties are in some further condition such that metaphysical faculties produced by unguided processes of the sort described by contemporary evolutionary theory and in that further condition are likely to be reliable—then by our Principle of Reason, the rational attitude for you to take toward the deliverances of your metaphysical faculties is doubt: agnosticism or disbelief. Assuming you have no good reason, as we'll put it, to doubt the doubts I have raised, then, if you are a practitioner of abstract naturalistic metaphysics, the rational course for you is systematic doubt toward the results of your work. And that's a deep problem for naturalistic metaphysics, since presumably the point of the enterprise is to arrive at claims one can believe, or at any rate, accept with considerable confidence. It's a problem, I'd think, for Flanagan,

Sarkissian, and Wong, since, one suspects, their attitude toward their various claims about metaphysical freedom, agent causation, the non-existence of a faculty of the will, there being no irreducible, non-natural moral properties, all moral truths being grounded in facts about human needs, desires and purposes, and so forth, is neither agnosticism nor disbelief. If they have no good reason to doubt the above doubts, however, the rational response to these various claims of theirs is either agnosticism or disbelief.

The big question, then: Do practitioners of naturalistic metaphysics have good reason to doubt the doubts I have raised about the reliability of those parts of our cognitive endowment responsible for abstract metaphysical beliefs? I don't think so, but let us look into some possibilities.

Reasons to doubt the above doubts?

A "miracle" argument?

There is an argument in the philosophy of science literature for scientific realism, the view that our scientific theories are true or approximately true (as opposed to just useful predictive devices), called the "miracle" argument, which might be relevant here. It takes its start from the observation that our best scientific theories are extraordinarily successful: they enable remarkably accurate predictions of a huge variety of phenomena, and enable us to manipulate and control our environment in myriad ways. The argument suggests next that by far the best explanation of this is that our theories are true or at least approximately true; for if they weren't, it would be nothing short of a miracle that they have been so successful. The most reasonable thing to think, the argument concludes, is that our theories are true or approximately true.

So far, the miracle argument for scientific realism. It has obvious application to my above claim that we have no reason to doubt the doubts I raise about the reliability of our metaphysical faculties. For if it's right, and the miracle argument provides good reason for accepting scientific realism, then doesn't the miracle argument likewise provide good reason for thinking our metaphysical faculties are reliable? Our best scientific theories involve quite a lot of abstract metaphysical theorizing (think relativity, and quantum physics), and if they are true or near true, then our metaphysical faculties would seem to be reliable.

To make the same point differently, suppose our metaphysical faculties are unreliable. How to explain, then,

our remarkable success at predicting and controlling the huge variety of phenomena we manage to predict and control by way of scientific theorizing? By far the best explanation of this, you might argue, is that our metaphysical faculties are reliable; for if they weren't, it'd be nothing short of miraculous that our scientific theorizing has been so successful. So the reasonable thing to think here, you might conclude, is that our metaphysical faculties are reliable and that we *do* have good reason to doubt the doubts raised above about their reliability.

Note by way of reply, though, that there are other ways of explaining the success of scientific theories, ways incompatible with scientific realism: e.g., Bas van Fraassen's constructive empiricism (1980). On van Fraassen's view, theory change in science is a sort of evolutionary process (39ff), a process that selects for theories that are *empirically adequate* (accurate in predicting and describing the observable world), weeds out theories that aren't empirically adequate, but is blind to what we might call *metaphysical adequacy*: accuracy in mirroring the unobservable structure of things. It's no miracle that an evolutionary process of this sort should produce successful (i.e., empirically adequately) but false (i.e., metaphysically inadequate) theories: you'd expect as much from such a process.

Now, it's no part of my project to argue that van Fraassen's constructive empiricism is correct. I bring it up only to point out that you'd have good reason for thinking the reliability of our metaphysical faculties the best explanation of our theoretical successes only if you had good reason for ruling out alternative explanations of the relevant successes, like van Fraassen's. The trouble is, there is furious debate in the philosophy of science over whether these alternatives are true, and the issues relevant to deciding the matter are quite, well, abstract. Deciding whether scientific realism or an alternative (non-realistic) explanation is the best explanation of our theoretical success, then, is itself an exercise in abstract metaphysics. Since deciding between these competing explanations is a crucial step in the miracle argument, we see that the miracle argument itself is a fine example of abstract metaphysics.

But no such argument could, all by itself anyway, constitute good reason for thinking your metaphysical faculties reliable. Suppose the defense attorney proposes to convince the jury that so-and-so is a reliable witness. The attorney's only evidence: the witness in question *says* of himself that he is reliable. That's pretty weak evidence; that a witness assures us he is a reliable witness isn't much reason

for thinking so. Likewise here: that you have an abstract metaphysical argument whose conclusion is that your metaphysical faculties are reliable isn't much reason for thinking so. Such an argument presupposes the reliability of the very faculties it purports to show the reliability of.[7] You couldn't sensibly accept the premises and method of reasoning in such an argument unless you *already* had some justification for thinking your metaphysical faculties reliable, justification which was prior to and independent of this argument. If your *only* evidence for thinking your metaphysical faculties reliable were this argument, it would be weak evidence indeed.

I conclude that if the miracle argument is your only reason for doubting the doubts I raised above about the reliability of your metaphysical faculties, then you have no good reason for doubting those doubts.

We should expect metaphysical reliability from creatures capable of abductive reasoning

Trent Daugherty offers another reason for doubting those doubts (Dougherty 2012; I'll put the objection in my own way, but the basic idea is his). Consider the fact that we have good reason to think that our Pleistocene forbearers were capable of abductive reasoning: reasoning to hypotheses about unobserved phenomena as the best explanation of observed phenomena. (This comes to us from anthropological studies of primitive hunter-gatherer societies, whose members are quite similar in terms of culture and cognitive power to humans living at the end of the Pleistocene era, and who make use of sophisticated abductive reasoning in hunting and tracking animals [see, e.g., Caruthers 2004].) Somehow or other, our ancestors evolved that capacity. But once they had that capacity, it's not surprising that they would have been capable of reliably redeploying it to previously unencountered domains. Think about arithmetic: once you have acquired the ability to reliably count buffalo, it's not terribly surprising that you would then be capable of reliably counting tigers. Likewise with deductive reasoning. Once you have acquired the ability to validly infer truths in some one domain, it's no great surprise that you would thereby be capable of valid inference in other domains as well.

Likewise with abductive reasoning. Our ancestors evolved faculties capable of reliable inference to the best explanation in the domains they encountered on the plains of Pleistocene Africa. That ability was adaptive and so passed

along to descendants. We have simply re-applied that method of reasoning to other domains, including the domains studied by abstract metaphysics, and it's no more surprising or improbable that it should be reliable in these other domains than it is that arithmetic or deductive reasoning should be reliable across a variety of domains.

We can put the point like this: True enough, we have good reason for thinking our metaphysical faculties are the product of naturalistic evolution and good reason for agnosticism about the probability that they should be reliable given that they are the product of naturalistic evolution. But we also have good reason for thinking our ancestors developed the ability to reason abductively and for thinking it likely that creatures capable of abductive reasoning should be capable of reliable metaphysical reasoning. The latter constitutes good reason for doubting the doubts raised above about the reliability of our metaphysical faculties.

So says our objector. By way of evaluation, let us look into the nature of abductive reasoning more closely. [8] Abductive reasoning, once again, is reasoning to hypotheses about unobserved phenomena as the best explanation of some observed phenomena. You emerge from a night's sleep to discover a bowl and spoon covered with dried ice cream on the counter. There is a variety of possible explanations: your house was broken into and the intruder enjoyed a bowl of ice cream but left all else untouched; your in-laws visited last night, had a bowl of ice cream, then left; and many more. But the best explanation is that your child had a late night snack before bed. Since that's the best explanation of your observational evidence, that's the one you believe.

Note two crucial features of the abductive process. First, there is a creative component: one imaginatively constructs possible explanations, where often these involve unobserved or unobservable mechanisms. In the reasoning of the previous paragraph, this part goes quickly, and you likely don't do much by way of imaginative construction of possibilities. Given your evidence, you likely need only consider a couple possibilities (or even just one). In other cases, in the sciences, for example, this part of the process is considerably more involved: one spends enormous time and effort imaginatively constructing various explanatory mechanisms. The key point: there is a creative moment in abductive reasoning, involving imaginative construction of explanatory mechanisms.

Second, there is a set of selection principles guiding one's choice of explanations (by which one determines the

best explanation). There is debate among philosophers and cognitive scientists about just what these are, but a typical list would include features like simplicity (the explanation is capable of simple expression, and postulates fewer entities or types of entities than its competitors), fit with background evidence (the explanation fits with other things one knows about the world better than competitors), explanatory scope (the explanation explains a wider range of data than do its competitors), and fruitfulness (it suggests new avenues of inquiry). The best explanation is the one that bests its competitors with respect to these and like selection principles.

Reflection on these two aspects of abductive reasoning suggests two requirements on successful use of abductive reasoning. First, there is what we might call the creativity requirement: successful abductive reasoning in some domain requires that we be capable of imaginatively constructing explanatory mechanisms of the sorts operative in that domain. Conceivably, a domain of inquiry could be so foreign to us, so beyond our ken, that construction of explanatory mechanisms of the sorts operative in that domain is beyond our imaginative powers. Successful abductive reasoning in a domain requires that that not be the case.

And second, there is what we might call the selection requirement: successful use of abductive reasoning in some domain requires that the selection principles guiding our adjudication between competing theories (simplicity, fruitfulness, etc.) are good indicators of the sorts of explanatory mechanisms likely to be operative in that domain. Conceivably, a domain of inquiry could be such that our usual principles of theory selection aren't good indicators of the sorts of mechanisms operative in that domain. Successful abductive reasoning in some domain requires that that not be the case.

These reflections on abductive reasoning in hand, return to this claim from above:

> Our ancestors evolved faculties capable of reliable inference to the best explanation in the domain they encountered on the plains of Pleistocene Africa. That ability was adaptive and so passed along to descendants. We have re-applied that method of reasoning to other domains, and it's no more surprising or improbable that it should be reliable in these other domains, including the domains studied by abstract metaphysics, than it is that arithmetic or

deductive reasoning should be reliable across a variety of domains.

We can now see that whether it is surprising or improbable that the abductive abilities we inherited from our Pleistocene ancestors are reliably redeployable in other domains, including those studied by abstract metaphysics, turns on whether it is surprising or improbable that our abductive powers satisfy the creativity and selection requirements vis-a-vis those domains. The crucial questions, then: How probable is it that the imaginative powers we inherited from our Pleistocene ancestors would have rendered us capable of imagining explanatory mechanisms of the sorts operative in the domains studied by abstract metaphysics (e.g., the quantum realm, or the realm of abstract objects)? And: How probable is it that the selection principles inherited from our Pleistocene ancestors [9] reliably discriminate between explanatory mechanisms likely to be operative in the domains studied by abstract metaphysics and those not likely to be operative in those domains?

By way of initial answer to these questions, I'd say: No one knows; no one is in a position to answer these questions. The main explanation why our Pleistocene ancestors evolved the imaginative powers and abductive principles they did was that having those powers, and making use of those principles, was adaptive in Pleistocene environments: useful for feeding, flying, fighting, and reproducing on the plains of Pleistocene Africa. Now, *maybe* imaginative powers and abductive principles selected for their success at those tasks in those environments would be capable of redeployment to the sorts of domains studied by abstract metaphysics, but it's hardly clear that we should *expect* this. Just as plausibly, the imaginative powers and abductive principles evolved by our Pleistocene ancestors were adaptive vis-a-vis evolutionary problems faced by hunter-gatherers on the plains of Africa but utterly unsuited to the domains studied by abstract metaphysics. As best I can tell, the probabilities in question are inscrutable: no one knows or could know what they are. We don't know how probable it is that abductive abilities evolved by our Pleistocene ancestors should satisfy the creativity and selection requirements vis-a-vis domains studied by abstract metaphysics: the quantum domain, etc. And if so, though it might be right that we have good reason for thinking our ancestors evolved the ability to reason abductively, we don't have good reason for thinking it likely that creatures with such cognitive capability should also be capable of reliable metaphysical reasoning. We don't, so far

forth, have good reason for doubting the doubts I have been on about.

A line of objection worth briefly exploring goes like this: "You are presupposing that the imaginative powers and abductive principles we modern humans deploy in our metaphysical theorizing are the same powers and principles our Pleistocene ancestors deployed in their theorizing. But this assumption is unwarranted. There is every reason to believe that these powers have evolved over the thousands of generations that separate us from our Pleistocene forbearers. And if so, it's no reason to doubt the reliability of *our* abductive practice in abstract metaphysical domains that *their* powers of imagination and abductive selection principles were inadequate to metaphysical theorizing."

There is a debate in cognitive science that is relevant here. Both sides of the debate accept a computational view of the mind, according to which the brain is an extremely complex digital computer and thought is a kind of computation. Both sides hold that the *hardware* of the brains of our Pleistocene forbearers is essentially the same as the hardware of our modern brains. The innate structure and consequent computational capacities of their brains differed very little from the innate structure and consequent computational capacities of our brains. (There was too little time between them and us for substantial hardware evolution.)

There is disagreement, however, about *software* differences between us and our forbearers.[10] One side of the debate holds that the software running on our brains has been thoroughly reprogrammed since the time of our Pleistocene forbearers. According to Daniel Dennett (Dennett, 1993), for example, software is encoded in human brains by language use and imitation, especially in childhood. The contents and information processing algorithms encoded in this software are determined by one's cultural environment, so that as culture evolves over time (new things learned about how the world works, new methods of reasoning discovered), so do the contents and information processing patterns programmed into young human minds. Cultural evolution, then, brings in its train software evolution, with the result that our brains are running software of an extremely different sort than that of our distant, Pleistocene ancestors.

Another side of the debate holds that there are relatively few software differences between us and our Pleistocene ancestors.[11] Their software enabled them to perform a host of information processing tasks which were

useful in solving the various evolutionary problems they encountered. We inherited from them the same basic suite of software, with those same basic information processing strategies. Proponents of this "continuity view"[12] of our mental software differ over the nature of the information processing strategies we inherited. Some think our information processing abilities "domain specific," so that our cognitive system is a collection of specialized faculties responsible for producing belief about particular domains. For example, some think we have a specialized belief-producing faculty responsible for producing beliefs about the mental states of others on the basis of observations of their face or body language, and a distinct, specialized belief-producing faculty responsible for producing beliefs about the behavior of objects in motion, and a distinct faculty yet for beliefs about which objects are living and which are mere artifacts, and many more (e.g., Pinker 1997). Others hold that our information processing abilities are "domain general," so that we have a collection of belief-forming strategies which apply generally, across the variety of domains we encounter: e.g., deductive logic, and the logic of probability (e.g., Gopnik and Melzoff 1997). But the key point about the continuity view, for our purposes, is that there are relatively few software differences between us and our Pleistocene ancestors: their software enabled them to perform a host of information processing tasks which were useful in solving the various evolutionary problems they encountered, and we inherited from them that same basic package of software with those same basic information processing strategies.

The continuity/no-continuity debate bears on our objection in obvious ways. If the continuity side is right, there likely isn't much difference between the imaginative powers and abductive principles deployed by our Pleistocene ancestors and those deployed today, and the points I raise above about our being properly agnostic about the adequacy of those imaginative powers and abductive principles to abstract metaphysics goes through as before.

Suppose the no-continuity side is right, though. What then? Might it be that the imaginative powers and abductive principles developed in us since our Pleistocene days have developed in the direction of metaphysical reliability?

Of course it might, but I think we have no way of knowing: that we have no way of knowing whether the software upgrades from our Pleistocene predecessors conduced to abductive reliability in metaphysical domains. How could we know? It's not self-evident or by any means

obvious that the software upgrades delivered to us by cultural evolution should have conduced to reliability in metaphysical domains. Perhaps the brain hardware evolved by our Pleistocene ancestors cannot be programmed by language-based instruction and imitation so as to subtend imaginative powers and abductive principles requisite for reliability in the quantum domain and other domains of abstruse metaphysical inquiry. Just as, for example, chimpanzee brains can't be so programmed: so it would seem anyway, no regime of language-based instruction or imitation can encode software in their brains sufficient to make them good at metaphysical theorizing. Perhaps the hardware of our brains is similarly constrained. It looks to me like the only reason one could have for thinking the hardware of our brains not so constrained, and for thinking that the software upgrades we inherited tended in the direction of metaphysical reliability, is some sort of abductive argument: argument, first, that the no-continuity view is the best explanation of this or that range of phenomena, and secondly, that the conjunction of the no-continuity view and the claims that (a) the brain hardware evolved by our Pleistocene forbearers was capable of software upgrades of the sort requisite for metaphysical reliability, and (b) in fact such upgrades occurred, is the best explanation of some further range of phenomena.

But any such argument would itself constitute an exercise in abstract metaphysics. It'd involve quite a lot of postulation of structure far removed from the world of everyday appearance, related to the world of appearances in highly complex ways, and so forth. And for reasons we saw earlier, that you have an abstract metaphysical abduction whose conclusion is that your abductive faculties are reliable in metaphysical domains isn't much of a reason for thinking them reliable. Such an argument presupposes the reliability of the very faculties it purports to show the reliability of. You couldn't sensibly accept the premises and method of reasoning in such an argument unless you *already* had some justification for thinking your abductive faculties reliable in metaphysical domains, justification which was prior to and independent of this argument. It's hard to see, though, what sort of justification that could be.

I conclude that, even if it's right that our imaginative powers and abductive principles have changed over time, we have no good reason for thinking they changed in the direction of metaphysical reliability, and thus no good reason here for doubting my doubts about the reliability of our metaphysical faculties.[13]

Help from Thomas Reid?

The 18th-century Scottish philosopher Thomas Reid famously argued that belief in the reliability of our cognitive faculties is not plausibly thought of as based on argument or inference from other beliefs we hold. According to Reid,[14] belief in the reliability of our faculties is based, rather, on a certain sort of *experiential* evidence, viz. an experience of the emotion of *ridicule*. He thought of belief in the basic truths of logic, arithmetic, and other truths of "common sense" as grounded in this emotion, as follows: When one considers the contrary of such truths—when one entertains the possibility, for example, that 1+1 is *not* equal to 2—an emotion of ridicule naturally arises, prompting in one thoughts like "that's absurd, ridiculous!". Such experiences, thought Reid, ground—make reasonable—belief in those truths of common sense whose contraries present themselves to one as thus ridiculous. In the language I've been using: such experiences constitute good reason for accepting such truths.

Perhaps we have such reason to doubt the doubts I raised earlier about the reliability of our metaphysical faculties. Perhaps belief in their reliability is like belief in the truths of common sense discussed by Reid: such that, when you consider the possibility that the contrary is true, that your metaphysical faculties are *not* reliable, you find yourself in the grip of a powerful emotion of ridicule. If so, we could argue à la Reid that you have good reason for thinking those faculties reliable, and good reason for doubting the doubts I raised above about their reliability. The reason doesn't take the form of an *argument* that they're reliable: it's a non-propositional, experiential reason. But no matter: it's good reason all the same.

By way of reply, though, I wonder this. Why don't *I* have this Reidian experiential evidence for the reliability of *my* metaphysical faculties? I don't experience any emotion of ridicule when I entertain the possibility that my cognitive faculties aren't reliable with respect to abstruse matters of metaphysics far removed from the everyday concerns of life. That possibility doesn't strike me as ridiculous at all. In fact, when I consider the multitude of crazy metaphysical views philosophers have defended over the centuries and the rampant disagreement among theorists over almost everything of substance on questions of metaphysics, I find it wholly *unobvious* that we humans, myself included, have reliable metaphysical faculties. More, that attitude strikes me

as *appropriate*, as the right response to the evidence. Given the history of crazy views defended over the years and rampant disagreement among practitioners of metaphysics, one *shouldn't* find ridiculous the possibility that our metaphysical faculties are not reliable—one shouldn't find it just *obvious* that our metaphysical faculties are reliable. Those who do find it thus obvious aren't, I think, being appropriately responsive to the evidence. But if so, then if your only evidence for belief in the reliability of your metaphysical faculties is Reidian experiential evidence, you *don't* have good reason to doubt the doubts I raised above about their reliability.

Conclusion

I have argued that commitment to naturalism gives good reason to doubt the reliability of those cognitive faculties responsible for abstract metaphysical theorizing. I claimed that, if there aren't good reasons to doubt these doubts, then commitment to naturalism gives one good reason for agnosticism about the deliverances of those faculties. I canvassed some main reasons for doubting these doubts. None were impressive. Since I think there aren't any better reasons, I conclude that there is trouble for the project of attempting to explain the mental, the moral, and the physical in naturalistic terms. Any such project will be an exercise in abstract metaphysics, suggesting that, if you are a participant in that project, the rational course for you is systematic doubt toward the results of your work.

[†]Thanks to Nathan Ballantyne, Justin Barrett, Matthew Braddock, Kelly Clark, and Gregg Ten Elshof for helpful comments and conversation.

[1] It's worth noting that naturalism, as I have defined it, is not the same thing as atheism. Naturalism, on my definition, says there is no good reason to believe in God, but it could be, of course, that though we have no good reason for thinking so, there nevertheless is such a being as God. So naturalism doesn't straightforwardly entail atheism. Nor does atheism straightforwardly entail naturalism. You could be an atheist, in that you deny the existence of an all-powerful, all-knowing, wholly good creator, but join Hegel, for instance, in thinking that total explanation of the world requires appeal to a cosmic mind of a non-theistic sort. The key point, then: naturalism, on my definition, is not the same view as atheism.

[2] The next two paragraphs owe an obvious debt to Peter van Inwagen's (2009,vi-xi) characterization of metaphysics.

[3] See, e.g., Plantinga (2000) and Beilby (2002).

[4] My argument also bears close affinities to arguments developed in two earlier papers of mine (2011; 2013), though below I'll develop a couple

new objections to the argument of those papers. Argument in the same basic family may also be found in Otte (unpublished), Reppert (2003), and Nagel (2012).

[5] See, e.g., Fales (2002), Fitelson and Sober (1998), Fodor (2002) and Ramsey (2002).

[6] I shall understand talk of *disbelief* here and in the sequel in such a way that someone disbelieves a proposition p if and only if she either fully believes not-p or invests a considerable degree of confidence in not-p. I understand talk of *agnosticism* here and in what follows in such a way that agnosticism regarding a proposition p is incompatible with considerable confidence in either p or not-p.

[7] Cf. Alston (1991, Chapter 3).

[8] I was helped here by Peter Caruthers' (2004, 81ff) discussion of abductive reasoning.

[9]For argument that those principles *were* inherited, see Caruthers (2004, 92ff).

[10] My discussion here is heavily indebted to Carruthers (2004).

[11] So says the "evolutionary psychology" approach of Leo Cosmides and John Tooby (see, e.g., Barkow, Cosmides and Tooby [1992]), Steven Pinker (1997), and collaborators.

[12] I borrow this terminology from Carruthers (2004).

[13] Similar considerations apply to the possibility, considered above, that reliability on abstract metaphysical matters is a "spandrel"—a non-adaptive byproduct of an adaptively selected trait. I said there that the prospects for producing an argument that reliability on abstract metaphysics is a spandrel are dim. We can now see why. Any such argument would constitute both a piece of abstract metaphysics and an argument for the reliability of our metaphysical faculties, and for reasons we have been considering, such arguments offer little by way of independent support for their conclusion.

[14] See, for example, Reid (1969, 606-607).

References

Alston, William P. 1991. *Perceiving God: The Epistemology of Religious Experience*. Ithaca, NY: Cornell University Press.

Barkow, Jerome, Leo Cosmides and John Tooby. 1992. *The Adapted Mind*. Oxford: Oxford University Press.

Beilby, James K. 2002. *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. Ithaca, N.Y.: Cornell University Press.

Carruthers, Peter. 2004. "The roots of scientific reasoning: infancy, modularity and the art of tracking." In Peter Carruthers, Stephen Stich, and Michael Siegal, eds., *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.

Crisp, Thomas M. 2011. "An Evolutionary Objection to the Argument from Evil." In Kelly James Clark and Raymond Van Arragon, eds., *Evidence and Religious Belief*. Oxford: Oxford University Press.

Crisp, Thomas M. 2013. "On coercion, love, and horrors" *Religious Studies,* 49: 165-179.

Dawkins, Richard. 1986. *The Blind Watchmaker*. New York: WW Norton.

Dennett, Daniel. 1991. *Consciousness Explained*. Penguin Press.

Dougherty, Trent. 2012. Review of *Evidence and Religious Belief*. *Notre Dame Philosophical Reviews*. http://ndpr.nd.edu/news/28918-evidence-and-religious-belief.

Fales, Evan. 2002. "Darwin's Doubt, Calvin's Calvary." In James K. Beilby, ed., *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. Ithaca, N.Y.: Cornell University Press.

Fitelson, Brandon and Elliot Sober. 1998. "Plantinga's Probability Arguments Against Evolutionary Naturalism," *Pacific Philosophical Quarterly* 79:2, 115–129.

Fodor, Jerry. 2002. "Is Science Biologically Possible?" In James K. Beilby, ed., *Naturalism Defeated? Essays on*

Plantinga's Evolutionary Argument Against Naturalism. Ithaca, N.Y.: Cornell University Press.

Gopnik, Andrew and Alison Melzoff. 1997. *Words, Thoughts and Theories*. Cambridge: MIT Press.

Nagel, Thomas. 2012. *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature is Almost Certainly False*. Oxford: Oxford University Press.

Otte, Richard. Unpublished. "Science, Naturalism, and Self-Defeat."

Pinker, Steven. 1997. *How the Mind Works*. Allen Lane.

Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.

Plantinga, Alvin. 2002. "Introduction." In James K. Beilby, ed., *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. Ithaca, N.Y.: Cornell University Press.

Ramsey, William. 2002. "Naturalism Defended." In James K. Beilby, ed., *Naturalism Defeated? Essays on Plantinga's Evolutionary Argument Against Naturalism*. Ithaca, N.Y.: Cornell University Press.

Reid, Thomas. 1969. *Essays on the Intellectual Powers of Man*. Cambridge: MIT Press.

Reppert, Victor. 2003. *C.S. Lewis's Dangerous Idea: In Defense of the Argument from Reason*. Downer's Grove, IL: InterVarsity Press.

Van Inwagen, Peter. 2009. *Metaphysics*. Boulder, CO: Westview Press.

van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Oxford University Press.