The Argument from Moral Responsibility

John Maier

Forthcoming in the *Australasian Journal of Philosophy*

Abstract

There is a familiar argument for the falsity of determinism, an argument that proceeds from the claim that agents are morally responsible. A number of authors have challenged the soundness of this argument. I pose a different challenge, one that grants its soundness. The challenge is that, given certain plausible assumptions, one cannot know the conclusion of this argument on the basis of knowing its premises. That is, one cannot know that determinism is false on the basis of this argument *even if* agents are in fact morally responsible and moral responsibility is in fact incompatible with determinism. A slightly different version of the challenge tells also against the claim that one can be justified in believing that determinism is false on the basis of the argument, so that the challenge cannot be evaded by a retreat to an epistemic position weaker than knowledge. I compare my challenge to the challenge posed by the external world skeptic, and argue that there are asymmetries between these challenges that make it reasonable to accept the former and reject the latter. I close by considering the prospects for developing an epistemology of moral responsibility that is adequate to answer the challenge.

1. The argument stated

The following argument, or one like it, has attracted a fair amount of philosophical attention:

(P1) Agents are morally responsible

(P2) Agents are morally responsible only if determinism is false

(C) Determinism is false

Call this the *argument from moral responsibility*, and call someone who claims to know (C) on the basis of this argument a *libertarian*.

The reason why this argument has attracted such attention is simply stated. Its first premise is intuitively compelling. So is its second premise – or, at least, it can be made so by some further argument.[1] And the argument is clearly valid. So we appear to have a sound argument for the falsity of determinism.[2] Yet the libertarian position is initially absurd. The falsehood of determinism, one would have thought, is something that one can know, if one can know it at all, only on the basis of considerations drawn from physics. It is not the sort of thing that can be known merely by reflecting on moral responsibility and a priori considerations. Yet that appears to be precisely what the libertarian has, with this argument, done.

Accordingly, a number of philosophers have attempted to show that this argument is not in fact sound. The *nihilist* denies the first premise: she says that agents are never morally responsible. The *compatibilist* denies the second: she says that agents may be morally responsible even if determinism is in fact true. Much has been said on behalf of

---

[1] There are at least two ways of giving the further argument. One might, in the manner of Peter van Inwagen [1983] and others, argue that determinism would rob agents of alternative possibilities, and that alternative possibilities are required for moral responsibility. Or one might, in the manner of Robert Kane [1996] and others, argue that determinism would prevent agents from being the source of their actions, and that sourcehood is required for moral responsibility. For the points to be made in what follows, it does not matter in which of these ways, if either, one argues for (P2).

[2] Determinism is here understood as the claim that a complete description of the world at any given time (say, a century before a certain purportedly morally responsible agent was born) nomically necessitates a complete description of the world at any other time (say, the time at which agent commits an act for which she is purportedly morally responsible). Nothing will hang on this particular way of formulating the claim: what is crucial will be the fact that determinism is (in some sense) an 'empirical' claim. It is an interesting question to what degree the arguments of this essay might generalize to other candidate empirical conditions on moral responsibility.

and against these denials.

I myself am unsure what to think. In certain moods nihilism seems plausible, and in certain moods compatibilism does. (Indeed, sometimes both can seem plausible: there is no moral responsibility, but even if there were, it would be compatible with determinism). But in other moods neither of these denials seems tenable. Yet the libertarian position remains unacceptable: one cannot come to know the falsehood of determinism on the basis of such considerations.

This essay should be of interest to those who are more confident in these matters than I am. But it is addressed in the first place to those who share my doubts. Its question is: how may one reasonably resist the libertarian position, without endorsing either nihilism or compatibilism about moral responsibility?

At first pass the pursuit of this question may seem to be hopeless. There would appear to be only three reasonable options here: to endorse nihilism, or compatibilism, or libertarianism. But matters may not be so simple. For note that the question of the soundness of the argument from moral responsibility is a question of logic (insofar as it is concerned with the validity of the argument) and of metaphysics (insofar as it is concerned with the truth of the premises). Yet the libertarian is advancing an epistemic claim: she represents herself as knowing a proposition (namely (C)), and furthermore as knowing it on certain grounds (namely (P1) and (P2)). If there is room for resisting each of the three allegedly exhaustive positions, it will be found in the gap between logical or metaphysical considerations and distinctively epistemic ones.

That there may be such a gap can be made clear by reflection on arguments that, though they may be sound, do nothing to support the position that is their conclusion.

Familiarly, the following argument is valid and may even, some think, be sound:

(P1') God exists

(C') God exists

Yet it would be unhappy for a theist to claim to know (C) on the basis of this argument. For if one is unsure about (C), this argument ought not change one's mind. Of course, the flaw in this argument is clear. It does not satisfy one of the conditions on good arguments once identified by G.E. Moore [1939]: its conclusion does not differ from its premise. If there is a flaw in the argument from moral responsibility, it must be more subtle.

I will argue that the flaw with the argument from moral responsibility is another of those identified by Moore. Moore claimed that for an argument to be a good one, its proponent must know its premises [1939: 166]. This is a standard that the proponent of the argument from moral responsibility (that is, the libertarian) cannot meet.

The argument for this position will be an extended one, but the basic idea is simple. Given certain plausible assumptions, the alleged epistemic position of the libertarian is an impossible one. This is because if someone is otherwise unsure about (C), then, if she knows (P2), (P1) is unknowable for her. So the libertarian is not in a position to know both of the premises of her argument, even if it is in fact sound.

This has the air of a skeptical argument, and to a certain extent it is. But it is a local skeptical argument, one that does not turn on or generate a more general skepticism. This is for two reasons. First, the skepticism defended here is hypothetical rather than categorical: the claim is that if an agent knows a certain proposition, then she cannot know a certain other proposition. Second, though there are certain symmetries between my argument and familiar skeptical arguments, there are also asymmetries, ones which turn on

more basic asymmetries between the epistemology of moral responsibility and the epistemology of other domains. These asymmetries will be the subject of the third section.

The fourth and last section looks forward to what work remains to be done. The arguments of this essay show that on the most natural ways of thinking of the epistemology of moral responsibility, the libertarian position cannot be sustained. This last section asks whether any epistemology of moral responsibility could support the libertarian position, and what form such an epistemology might take. As the discussion there tries to make clear, this essay is less a refutation of libertarianism (whatever that would be) than it is an invitation for further development of the position, development on the epistemic rather than the much-explored metaphysical aspect of the position. This last section offers some tentative suggestions about how that invitation might be answered.

2. Why the argument fails

The case against the libertarian proceeds in three stages. First, I characterize her epistemic situation. Second, I defend three claims about someone who is in that situation. Third, I point out that those three claims entail that someone in the epistemic situation of the libertarian cannot know both (P1) and (P2); a fortiori, she cannot know (C) on the basis of knowing (P1) and (P2). A libertarian may then resist the argument in one of two ways: by denying that her epistemic situation has been adequately characterized in the first stage of the argument, or by denying one of the claims that figure in the second stage of the argument.

2.1 The epistemic situation of the libertarian

Whether an argument justifies its conclusion will depend in part upon the epistemic situation in which it is deployed. I have said that the libertarian argument fails when deployed by someone who is 'otherwise unsure' about its conclusion – that is, about the truth or falsity of determinism. What does this mean?

At first pass, we might say the following. An agent is otherwise unsure of (C) just in case she has no evidence for or against (C). But this cannot be right. For the libertarian herself claims to have evidence for (C) – namely her evidence, whatever it is, for (P1) and (P2). So this is not a construal of the libertarian's epistemic situation that could be accepted by the libertarian.

Somewhat more carefully, we might say the following. An agent is otherwise unsure of (C) just in case she has no evidence for or against (C) other than her evidence, whatever it is, for (P1) and (P2). That is, her evidence for the claim that agents are morally responsible and her evidence for the incompatibility of determinism and moral responsibility exhaust her evidence for (C).

This cannot be quite right either, however, if what is wanted is a sympathetic and accurate account of the libertarian position. For it will be pointed out that we ourselves have significant evidence for (C), such as that yielded by quantum mechanical considerations. This point involves some not insignificant assumptions about quantum mechanics and its proper interpretation. Nonetheless, let us grant it, and grant then that the libertarian has evidence for (C) aside from her evidence for (P1) and (P2). We then need to be somewhat more subtle in our characterization of the libertarian position.

To approach the characterization, note that though the libertarian may possess such evidence, this evidence seems in a certain way irrelevant to the epistemic issues that are at

stake here. We may bring this point out temporally. The argument from moral responsibility was deployed long before the discoveries of the last century's physics. Thus the temporal claim: the libertarian knows her conclusion now only if libertarians knew their conclusion then. Or, and somewhat more fundamentally, we may bring this out in terms of what a libertarian's epistemic position would be were she to lose her additional evidence. Thus the modal claim: a libertarian knows her conclusion only if she would know it even if she lost all of her other evidence for (C) (all of her evidence, that is, other than her evidence for (P1) and (P2)).

One might deny either the temporal or the modal claim; both of them raise interesting and rather deep issues about the nature of justification generally. Nonetheless, for present purposes, we can focus on our attention on the libertarian who is willing to accept them. For it is precisely such a libertarian who incurs the charge of absurdity made at the outset. What is remarkable about her position is that she purports to derive a physical conclusion from moral and a priori considerations alone. The epistemic viability of a 'mixed' libertarianism, one which depends upon both empirical and a priori considerations, is an interesting topic as well, but a topic for another essay.

Let us then introduce a notion of *effective ignorance* as follows. An agent is effectively ignorant of the conclusion of her argument just in case her argument would be exactly as epistemically effective as it actually is if she had no evidence for (C) other than her evidence for (P1) and (P2). In terms of effective ignorance, we can state more precisely the thesis that will be defended in the second stage of the argument. This is the following: someone who is effectively ignorant of (C) cannot know both (P1) and (P2). Thus the libertarian with whom we will be concerned cannot know both of the premises of her

argument, and so her argument fails, even if it is in fact sound. I will now give the case for that claim.

2.2 Particularity, authority, humility

The case will itself take the form of a deductive argument. Three claims, taken together, entail that if someone is in the epistemic situation just described with respect to (C), then, if she knows (P2), (P1) is unknowable for her. None of these three claims is undeniable, and one way of resisting the present argument is simply to deny them. Nonetheless, each of them is plausible, for reasons that I will explain in what follows.

Let us take the three claims in turn. All three claims will concern some arbitrary agent S who is effectively ignorant of (C) in the sense defined above. I will first state the claims, and then defend them.

The first claim is *particularity*. According to this claim, S can know (P1) only if she can know some instance of (P1).

The second claim is *authority*. According to this claim, S can know some instance of (P1) only if she can know some instance of (P1) on partly non-testimonial grounds. (I will explain what it is precisely for grounds to be 'partly non-testimonial').

The third claim is *humility*. According to this claim, if S knows (P2), then she cannot know any instance of (P1) on partly non-testimonial grounds.

These three claims together entail the *impossibility result*: if one is effectively ignorant of (C), then it is impossible to know both (P1) and (P2). Since one can come to know (C) on the basis of (P1) and (P2) only if one knows (P1) and (P2), it is therefore impossible to come to know (C) on the basis of (P1) and (P2). The libertarian's alleged

epistemic position turns out to be an impossible one.

Since the three claims clearly entail the impossibility result, the libertarian who accedes to the description of her epistemic situation given above and who wishes to retain her position must deny one of these claims.[3] Yet each of these three claims is supported by a plausible claim about the epistemology of moral responsibility, as I shall now argue.

Begin with particularity: S can know that agents are morally responsible only if she can know some instance of the claim that agents are morally responsible. Let us consider what it would be for some particular agent to fail to satisfy particularity, and then appeal to some more general considerations in virtue of which there seems to be something epistemically objectionable about such an agent.

Consider a psychiatrist, Jones, who is called on to offer testimony in a trial as to the responsibility of some defendants accused of burning down a factory. It is stipulated that the defendants have indeed burned down the factory. The question is whether they are responsible for having done so.[4] Jones holds that the defendants, taken ensemble, are indeed responsible for burning down the factory. Yet she claims to be ignorant of whether any individual defendant is responsible for his part in burning down the factory. We can imagine versions of the case in which this would be a reasonable position to hold. It may be, for instance, that the defendants constitute a 'mob' such that they are together

[3] Perhaps there is still another possibility, which is to grant the soundness of this argument without granting that one can know its conclusion on the basis of knowing its premises (that is, to adopt against my challenge to the argument from moral responsibility the strategy of the challenge itself). Since this kind of strategy is applicable in only a certain limited range of cases, as I will explain below, this does not seem an especially promising response, but nothing said here rules it out.

[4] What is at issue in these cases is legal responsibility, which is to be distinguished from the proper topic of this essay, namely moral responsibility. I elide the distinction here on the perhaps optimistic assumption that, at least when it comes to the epistemic issues that concern us here, legal and moral responsibility behave similarly. Those who do not grant this assumption may not be moved by the example of Jones, though they are still subject to the more general considerations that I will raise presently, considerations which apply directly to moral responsibility.

responsible for the burning without any individual being responsible for the burning. But let us say that Jones denies this. She holds simply that while the defendants are responsible, she is ignorant of whether any individual defendant is responsible. Here it is reasonable to ask Jones on what grounds she holds the defendants to be responsible, if not on the grounds that suffice to hold the individual defendants to be responsible as well. And if Jones does not have a response to this question, then we may conclude that she does not know that the defendants are responsible after all.

Particularity simply generalizes this criticism of Jones: someone's knowledge of moral responsibility must somehow or other involve knowledge of the responsibility of individual agents. Abstracting from the courtroom scenario, we can say something still more general in defense of particularity. It is an instance of a more general thought which is that justification in certain domains is 'constructive': one can know the general proposition only if one knows one of its instances. Color provides a natural example. It is plausible that (provided we are realists of some sort about color) we know that certain objects are colored. But this is not knowledge that we could arrive at independently of having knowledge of the particular colors of particular objects. Rather, this knowledge requires knowledge of the particular color of some particular object, for instance that this tomato is red. Particularity comes to the idea that moral responsibility is like color in this regard. Someone could not know merely that agents are morally responsible generally, and never know of some particular agent that she is morally responsible for some particular action (or omission). The epistemology of moral responsibility is in this respect particular.

This sort of constructive requirement on knowledge does not always hold, but considering the cases where it fails makes it more plausible that it does in fact hold of moral

responsibility. The most familiar case where the constructive requirement fails is perhaps the case of introducing posits in explanations. The physicist may infer that there are particles of a certain sort, without being able to know any instance of such a particle, because postulating such particles provides the best explanation of the phenomena. The reason why the epistemology of moral responsibility is particular is plausibly that moral responsibility does not play this sort of role for us: it is not introduced as something that explains the phenomena. Rather it is, if anything, one of the phenomena to be explained – here again it is like color. While there are certain exceptions to the constructive requirement on knowledge, these exceptions are importantly disanalogous to the case of moral responsibility. So reflection on the case of Jones and this more general consideration give us reason to accept particularity about moral responsibility: S can know that agents are morally responsible only if she can know some instance of the claim that agents are morally responsible.

Turn then to authority: S can know some instance of the claim that agents are morally responsible only if she can know some instance of the claim that agents are morally responsible on partly non-testimonial grounds. Let us consider what it would be for some particular agent to fail to satisfy authority, and then appeal to some more general considerations in virtue of which there seems to be something epistemically objectionable about such an agent.

First, however, we need to explain what it is for some knowledge to be 'partly non-testimonial.' It would be too much too demand that agents knowledge of moral responsibility not be based in testimony at all. Much of what we take to be knowledge of moral responsibility depends, in rather deep ways, on the testimony of others, be it

testimony from a religious or moral tradition, or from sophisticated a priori argumentation whose soundness an agent may well accept on the grounds of the testimony of others. And indeed in many cases our knowledge of moral responsibility may depend entirely on testimony of others. If a reliable source tells me that a certain person committed a robbery and was responsible for her action, then I may thereby come to know of that person that she was responsible for her action. To accept authority is to deny neither of these claims. Authority rather makes the very weak claim that an agent knows about moral responsibility generally only if she knows, in at least some case, on grounds that are at least in part non-testimonial, that some agent is morally responsible for what she has done. As such it is compatible with the thoroughgoing and perhaps necessary dependence of our knowledge of moral responsibility on the testimony of others.

Having thus clarified authority, let us try to evaluate it. Consider again Jones, now in a case where she claims to know that a given agent, Smith, is responsible for what he has done. Some of her knowledge of what Smith has done will of course be based on testimony. Yet if she is called on to explain why Smith himself is responsible, and she avers to someone else's testimony on this point, then we may reasonably ask on what grounds Jones herself knows Smith to be responsible. Again, we can imagine versions of the case in which Jones herself has an adequate answer to this question. She may for example point out that Smith's brain has developed in a particular way, and aver to some psychiatric authority who holds that agents whose brains have developed in this way are morally responsible. But let us imagine that Jones does not do this. She simply holds that Smith is responsible on the grounds that she has been told by someone else, Brown, that he is responsible. We may reasonably then deny that Jones knows Smith to be responsible, or

rather that she knows this only in virtue of this fact being known, in part non-testimonially, by Brown. And we should then want to ask Brown on what grounds he holds that Smith is responsible. And if Brown has no knowledge on this point that is not itself in part non-testimonial, then we will reasonably conclude that neither he nor Jones in fact knows that Smith is responsible for what he has done.

Note that, on a natural reading of this case, Jones and Brown do not in fact violate authority, strictly speaking. For authority demands only that someone have non-testimonial knowledge or other of moral responsibility, and this is presumably something that Jones and Brown do have, for they presumably have extensive non-testimonial knowledge of moral responsibility, concerning occasions other than the one under consideration. But we can alter this feature of the case by making the admittedly unnatural stipulation that Jones and Brown have no other non-testimonial knowledge concerning moral responsibility: that this is, as it were, their only purportedly direct encounter with taking an agent to be morally responsible for what he has done. This stipulation does not make the position of Jones and Brown any less criticizable; if anything, it makes it more criticizable.

Some may wonder whether the case we are being asked to suppose is even coherent. Note that, in order for Jones and Brown to have no non-testimonial knowledge of moral responsibility, they must not even have acquaintance with their own moral responsibility in the first-personal case. And perhaps one might doubt whether such a scenario can even be supposed. But here we must be careful about where we are in the dialectic. The claim we are trying to evaluate is authority: S can know some instance of the claim that agents are morally responsible only if she can know some instance of the claim that agents are morally responsible on partly non-testimonial grounds. What we are pointing out is that it

is at least very difficult to imagine cases where the consequent of this conditional is false: that is, it is very difficult to imagine cases where authority may be false. But, in that sense, authority imposes a very weak demand on creatures in the kind of epistemic situation that we take ourselves to be in. And since the claim at this point is is merely that it is true, the difficulty of conceiving counterexamples to it is, at this point at least, a virtue of authority.

Someone willing to grant the point just made may wish to raise a somewhat more subtle point about the role of these cases in the dialectic. According to this point, the case just made for the truth of authority threatens the case for the explanatory relevance of authority. Note that, in the unmodified version of the case, where we assume that Jones and Brown have some other non-testimonial knowledge of moral responsibility, we nonetheless take them to be criticizable. Can the present account explain this criticism, and does this criticism have anything to do with the condition called authority (if, as has just been noted, that condition is so easily satisfied)? I suspect that the answer to both questions is affirmative. This will be so on the supposition that, if authority is indeed a condition on knowledge of moral responsibility, a more circumscribed version obtains in a host of more ordinary cases: according to this circumscribed version, someone can know some instance of the claim that agents are morally responsible only if she can know some *relevant* instance of the claim that agents are morally responsible on partly non-testimonial grounds. In the original version of the case, Jones and Brown's non-testimonial knowledge of moral responsibility is not relevant in the required sense, and this is why we take them to be criticizable. A full defense of this diagnosis would require, among other things, an adequate account of relevance. If this could be done, then we would have an account of why Jones and Brown are criticizable in the initial case and why the criticism arises

ultimately from authority, for the grounds of this restricted version of the authority condition would ultimately be the authority condition itself.

Authority, then, simply generalizes these criticisms of Jones (and of Brown): someone's knowledge of the responsibility of some individual must somehow or other involve some partly non-testimonial knowledge of some individual's responsibility. Abstracting from the courtroom scenario, we can say something more general in defense of authority. While particularity was motivated by analogy with color, authority marks a disanalogy with color. Someone might know of various objects that they are colored though one's knowledge is based entirely on testimony. Someone who is congenitally blind, for instance, arguably might possess such knowledge. Authority claims that knowledge of moral responsibility is not like this. A natural defense of this claim appeals to a parallel claim defended in other areas of moral epistemology. Some philosophers, such as Alison Hills [Hills 2009], have argued that there is something odd about moral testimony: when my claim that (for instance) lying is always wrong bottoms out in deferring to some expert who informed me about this, there is something suspect about my purported epistemic scenario. The diagnosis of this phenomenon is not obvious, but it suggests at least that moral testimony all on its own fails (for some reason) to generate moral knowledge. And, if that is correct, it may be that it does not stop at the bounds of what is right and wrong. It is similarly odd to base my judgments about justice, for example, entirely on such grounds. Authority about moral responsibility simply extends this thought to moral responsibility. So reflection on the case of Jones and this more general consideration give us reason to accept authority about moral responsibility: S can know of any particular agent that she is morally responsible for her action only if she can know on partly non-testimonial grounds

that some particular agent is morally responsible.

Turn finally to humility: if S knows that agents are morally responsible only if determinism is false, then she cannot know any instance of the claim that agents are morally responsible on partly non-testimonial grounds. Let us consider what it would be for some particular agent to fail to satisfy something like humility, and then appeal to some more general considerations in virtue of which there seems to be something epistemically objectionable about such an agent.

Consider, once more, Jones. Now Jones is testifying, of Smith, that he is responsible for what he has done, and is not purporting to base her claim entirely on the testimony of others. There is something legally inapt about an appeal to determinism in such a case, but let us imagine that Jones holds that Smith is responsible only if it is false that he has a certain brain lesion. It is pointed out, by Smith's lawyer, that Jones is in no position to know whether Smith has that brain lesion. We can imagine versions of the case where Jones responds to this point by acknowledging that, on this point at least, she needs to defer to someone else. But let us say that she does not do this. She rather claims to know that agents who have a certain brain lesion are not responsible and also, though she is ignorant of whether Smith has this brain lesion, that Smith is responsible. Here we may reasonably protest that Jones does not know that Smith is responsible after all.

Humility simply generalizes parlays this criticism of Jones, with regard to the brain lesion, into a criticism of the libertarian, with regard to determinism. Abstracting once more from the courtroom scenario, we can say something more general in defense of humility. The illegitimacy of Jones's reasoning might be motivated by the following principle, which we might call the *epistemic priority of excuse*:

> If K is an excusing condition for some particular action A of some particular agent R, then one cannot know on partly non-testimonial grounds that R is morally responsible for A unless one independently knows that K does not obtain.

An 'excusing condition' is one such that if it obtains, then R is not morally responsible for A. Insanity, infancy, and so forth are typical examples of excusing conditions, as was the brain lesion in the case just described. So too, if (P2) is true, is determinism.[5] Therefore, if S knows (P2), the factivity of knowledge conjoined with the epistemic priority of excuse entail humility.

The epistemic priority of excuse should be regarded as a conjecture about the epistemology of moral responsibility, one that would systematize our intuitions about particular cases and also be the foundation for what we are calling the principle of humility. If it is true, it explains why it is typically not legitimate to infer that, for example, since someone is morally responsible, she is therefore not insane. According to the epistemic priority of excuse, this gets things the wrong way round. One must first figure out on independent grounds whether or not someone is insane, and then come to a judgment about whether she is morally responsible.

If this is right, then reflection on the case of Jones as well as these more general considerations give us reason to accept humility about moral responsibility: if S knows that agents are morally responsible only if determinism is false, then she cannot know any instance of the claim that agents are morally responsible on partly non-testimonial grounds.

What if, even after reflecting on cases of the sort described above, one simply does

---

[5] R. Jay Wallace [1994] argues that determinism is not an excusing condition under our ordinary conception of excuses. Perhaps he is right about this; certainly there are significant differences between infancy and insanity on the one hand and determinism is the other. There are also, however, similarities between these conditions, at least by the libertarian's lights, and this is why I gather them together under the term 'excusing condition,' the meaning of which is simply stipulated in the foregoing. The epistemic priority of excuse is an epistemic claim about moral responsibility and 'excusing conditions' when the latter is read in a wide, and perhaps technical, sense. As such it is compatible with the further distinctions that Wallace emphasizes.

not find the epistemic priority of excuse – which is, at this point, simply offered as a conjecture in the epistemology of moral responsibility – compelling? If this principle is a basic one in the epistemology of moral responsibility, one cannot give an argument from more fundamental premises for this claim, which on the present view constitutes something like an a priori principle about excuses. One can however pose the following challenge. Someone who rejects the epistemic priority of excuse either accepts humility, or she does not. If she does, then the question is on what grounds she accepts it, if not an appeal to the epistemic priority of excuse. If she does not, then the question is what on what grounds she resists the argument from moral responsibility. This challenge is hardly decisive: some compatibilists have an answer to the first question (see, for instance, the subsequent paragraph), while libertarians will answer the second question by saying that they do *not* want to resist the argument from moral responsibility. Here it is useful to emphasize the conditional nature of the argument: given certain simple claims, there is a novel way of resisting the argument from moral responsibility. If someone in the end simply does not find these claims compelling, then she may reasonably wish to endorse some other position in the free will disputes.

Even someone sympathetic with humility might find the line or reasoning just given unconvincing, or at least redundant. For it might seem that there is a quicker argument for humility, one which does not go via the epistemic priority of excuse. This argument, roughly put, is as follows. We clearly can have knowledge of instances of moral responsibility on partly non-testimonial grounds. If moral responsibility were the kind of thing that demanded the falsehood of determinism, then such knowledge would be impossible. So these epistemic considerations point us to a view of a less demanding view

of the metaphysics of moral responsibility: moral responsibility demands not 'the ability to do otherwise' but demands, at most, the presence of certain capacities, and these capacities are entirely compatible with the possible truth of determinism.

But given the present approach, the arguments for humility – which nowhere appealed to these sorts of claims about the metaphysical conditions on moral responsibility – are not at all redundant. For the view invoked in the previous paragraph is an instance, perhaps the most prominent contemporary instance, of a view set aside at the outset: this is the compatibilist view according to which the argument from moral responsibility is simply unsound in virtue of the falsehood of (P2). There are well-known objections to this compatibilist view, which is not to say that these objections are decisive.[6] For present purposes what needs to be emphasized is that this compatibilist view is quite different from the view proposed here, and that for those of us unpersuaded by the compatibilist way of resisting the argument from moral responsibility, the view described here represents an alternative to it.

2.3 The impossibility result

Particularity, authority, and humility jointly entail the impossibility result. Before coming to that result, it bears noting that the kind of challenge being posed here cannot be evaded simply by withdrawing the claim to know (C), for a parallel challenge tells also against an epistemically more modest position. Let a *modest libertarian* be someone who

[6] One set of objections concerns whether this kind of view can uphold all of our pretheoretical judgments about moral responsibility, especially those concerning omissions; such objections are discussed extensively by John Martin Ravizza and Mark Ravizza in the course of their defense of a version of this kind of compatibilism [1998: 123-150]. Another objection, closer to our present concerns, is more straightforwardly epistemic. In particular, Scott Sehon [2012] argues that the epistemology of moral responsibility may remain puzzling even on a capacity-based approach (Sehon's argument concerns Fischer and Ravizza's account specifically, but as he notes it might naturally be extended to other such approaches).

claims merely to be justified in believing (C) on the basis of (P1) and (P2), not to know it. And let S* be a modest libertarian who is effectively ignorant of determinism in the sense already given. The argument just given does not apply to S*. For the argument above relies, in its defense of humility, on the claim that S knows (P2). Since the libertarian claims to know (C) on the basis of knowing (P1) and (P2), this argument applies to her. But the modest libertarian need not make this claim. Since she claims only to be justified in believing (C) on the basis of (P1) and (P2), she can instead claim to be no more than justified in believing (P2). The argument for humility, and so the argument for the impossibility result more generally, does not therefore extend to her.

Yet a slight revision of the argument does allow for precisely such an extension. Let *modest humility* be the claim that if someone is justified in believing that agents are morally responsible only if determinism is false (that is, if she is justified in believing (P2)), then she cannot know any instance of the claim that agents are morally responsible on partly non-testimonial grounds. And let the *modest epistemic priority of excuse* be:

> If one is justified in believing that K is an excusing condition for some particular action A of some particular agent R, then one cannot know on partly non-testimonial grounds that R is morally responsible for A unless one independently knows that K does not obtain.

According to this claim, if I am justified in believing that (say) insanity is an excusing condition for a crime, and if I am effectively ignorant of whether a given agent is insane, then I cannot know on partly non-testimonial grounds that she is morally responsible for her crime (whether or not insanity is in fact an excusing condition for her crime). This principle seems at least no less secure than the epistemic priority of excuse itself, in the sense that anyone who endorses that principle ought also endorse this one, for this principle is just, as it were, a 'subjective' variation on the epistemic priority of excuse. Since the

modest libertarian is purportedly justified in believing that determinism is an excusing condition for moral responsibility, modest epistemic priority of excuse entails that modest humility is true of her, and this conjoined with the other claims entails a modified version of the impossibility result: the purported epistemic situation of S* is an impossible one. (Since the remarks developed below about the impossibility result apply also, mutatis mutandis, to the argument given in this paragraph, I shall set modest libertarianism to one side except where relevant).[7]

We have now arrived at the impossibility result. S can know (P1) only if she can know some instance of (P1). S can know some instance of (P1) only if she can know some instance of (P1) on partly non-testimonial grounds. And if S knows (P2) (and so if (P2) is true), S cannot know any instance of (P1) on partly non-testimonial grounds. From this it follows that S cannot know both (P1) and (P2). Given that one can know a conclusion on the basis of certain premises only if she knows those premises, we have shown that the purported epistemic position of the libertarian – according to which, effectively ignorant of (C), she comes to know (C) on the basis of (P1) and (P2) – is an impossible one.

There are two ways of resisting this argument, given that one accedes to the initial description of the libertarian's epistemic situation. First, one could deny one of the three claims that are its premises. I have given reasons in favor of these above, reasons which I

[7] A rather different variation on the libertarian position is someone who claims merely to be justified in believing that certain agents are morally responsible, and not to know this. Such a position evades the argument just given, and indeed the arguments this essay more generally. It faces a rather different sort of challenge. For this view of moral responsibility would seem to involve a fair degree of revision of our actual practice of attributing moral responsibility. For it seems we do purport to know, and not merely to be justified in believing, that agents are morally responsible for their actions or omissions on at least some occasions. This position salvages the argument from moral responsibility, then, only at the price of withdrawing that claim to knowledge. It may be that these sorts of revisions could be avoided or somehow made palatable, but the question of how and whether that sort of revision can be undertaken successfully is a quite different one from the questions that concern us here.

take to be sufficient, though of course there is more to be said about each of these. Second, and more indirectly, one could try to show that the reasoning employed here yields implausible conclusions when applied more broadly: that is, that the argument 'proves too much.' This might show that, even if we cannot pinpoint any particular error in the argument, we could know that there must be some error or other in it. The next section takes up this way of challenging the argument.

3. Why we have not proven too much

Something like this challenge is posed by Peter van Inwagen in the course of defending his own position on these issues.[8] With impressive frankness, van Inwagen [1983: 210] articulates, in an opponent's voice, the epistemic unease with which we began:

> If all of your arguments are correct, then our (alleged) knowledge of the existence of moral responsibility, coupled with certain arguments a priori, can constitute a good reason for believing that determinism is false. But these things are not the sorts of things that can be a good reason for believing in indeterminism. Indeterminism is, to put it bluntly, a thesis about the motion of particles of matter in the void . . . only scientific investigations are relevant to the truth or falsity of such theses.

van Inwagen's response is indirect. Rather than answering the challenge on its own terms, he argues [1983: 211] that anyone moved by the foregoing reasoning ought also be moved by the following reasoning, given in the voice of the skeptic:

> You say that most of the propositions we unreflectively assume we know to be true are true. You say, moreover, that we know this, or, at least, have good reason to believe it. But we can deduce from this thing you say we have good reason to believe that there exists no Cartesian Universal Deceiver – no being who deceives us all about almost everything . . . but these considerations are not the sorts of considerations that can provide us with good reasons for believing that there is no Universal Deceiver. The thesis that there is no Universal Deceiver is, to put it bluntly, a thesis about the features of a part of the world that is inaccessible to any

[8] van Inwagen's own position is a version of what we are calling modest libertarianism, since (as will become clear) it is put (partly) in terms of reasons for belief, rather than in terms of knowledge.

possible human investigation.

This reasoning, van Inwagen argues, is generally held to be objectionable, even if it is difficult to say just why.[9] If that is accepted, then the grounds for holding that reasoning to be objectionable carry over also as grounds for holding the reasoning deployed against his own position to be objectionable.

This section will, in effect, take up van Inwagen's challenge in a regimented way. For we have not merely expressed vague unease at the libertarian's epistemic situation. We have rather offered a clear argument that the libertarian's alleged epistemic position is an impossible one. If a similar argument can show that it is impossible to know that the external world exists, then perhaps we should accept that the foregoing argument proves too much (even if we can't say where exactly it goes awry). On the other hand, if we can find principled reasons for rejecting this skeptical argument that do not appear to apply to the argument given in the previous section, then we have a response to the challenge.

I will argue for the latter position. There are principled reasons for denying the analogues of both particularity and humility in the case of our knowledge of the external world. Indeed, these correspond to two traditional ways of responding to external world skepticism. So there are grounds for breaking the alleged symmetry between the argument against the libertarian and external world skepticism: one can consistently, indeed plausibly, accept the former and the reject the latter.

---

[9] van Inwagen considers two replies. One simply rejects the reasoning and says that one can after all have reason to believe that there is no Cartesian Universal Deceiver. The other denies the closure of having a reason to believe, and says that one can have reason to believe that (for instance) one has hands without thereby having reason to believe everything that is entailed by the claim that one has hands. van Inwagen says he favors the first response, but that for his purposes one could make the second as well ('This is not,' he remarks, 'a book about epistemology' [1983: 212]).

3.1 The denial of particularity

Begin with particularity. This, recall, was the claim that S can know (P1) only if she can know some instance of (P1). As noted above, the kind of entities that generally violate this requirement are the sorts of theoretical entities posited by the sciences. And one of the traditional ways of responding to external world skepticism is precisely to treat the objects represented in our experience as precisely, in a way, such posits. The thought is that our belief in external objects is justified by what we have come to recognize as a sort of inference to the best explanation, where the explanandum is the order and regularity of our sensations, and the hypothesis of external objects mind-independent that which explains this. Such an epistemology gives us an explicitly non-particular account, in the sense that it does not require of any particular external object that we know that it exists. Rather, external objects are on such a picture things which we know only indirectly, through their effects.

We said above that this was not how we conceived of moral responsibility. But the issue is not, or not merely, that this is not how we conceive of moral responsibility. More fundamentally, the issue is that moral responsibility it is part of a family of concepts – including but not exhausted by moral concepts – that appear to be explanatorily dispensable. One argument for this thought is given in an influential discussion by Gilbert Harman [1977: 6-9]. Harman notes that, confronting some hoodlums who set a cat on fire, we might wish to explain their behavior by citing the fact that it is wrong to do so. But Harman argues that the appeal to moral facts is not necessary. It is sufficient to appeal to beliefs that have moral concepts in their contents, in this case the hoodlums' belief that it is wrong to ignite the cat.

Similarly, it seems we need not invoke facts about moral responsibility in explanations of our attitudes and behavior. We need only invoke our beliefs about moral responsibility. The explanation for why we blame these children for their act, if we do blame them, is not that they are morally responsible. It is rather that we believe that they are morally responsible. And if moral responsibility has no place in our explanations of our attitudes and behavior, then it has no place in explanation simpliciter, for this is the explanatory home for moral responsibility, if it has any explanatory home at all. Thus the grounds for denying particularity in the case of external objects are not plausible grounds for denying it in the case of moral responsibility.

3.2 The denial of humility

Turn then to humility. This, recall, is the principle that if someone knows that agents are morally responsible only if determinism is false, then (since she is effectively ignorant of the truth of determinism) she cannot know on partly non-testimonial grounds of any particular agent that she is morally responsible for some particular action. An analogous principle in the case of external world skepticism would hold that if someone knows that she has hands only if she is not a handless brain in a vat then (since she is effectively ignorant of whether she is a brain in a vat) she cannot know on non-testimonial grounds that she has hands. The reasoning behind humility turned essentially on the thought that if one is ignorant of some proposition p, and that proposition is incompatible with a certain proposition q, then one cannot know q without having some grounds for ruling out p other than q itself. The legitimacy of such reasoning is not incontestable, and one of the main strands in recent discussions has been to contest it, at least to contest the thought that it can

play its required role in skeptical arguments.

There are a number of ways in which such a defense might be mounted, but a statement of the guiding idea behind them is given by James Pryor [2000: 536]:

> It seems like *the mere fact* that one has an experience of that phenomenological sort [viz. as of hands] is enough to make it reasonable to believe that there are hands. No *premises* about the character of one's experience – or other sophisticated assumptions – seem to be needed.

This is a natural thought, and if it is accepted then we have a natural way of resisting skepticism about the external world. The question that then arises is whether it can be extended to defending the libertarian position. It is a deep and seemingly open question how far this sorts of response may be deployed. There are at least two reasons, however, for thinking that it cannot be extended so far as to encompass moral responsibility.

The first is that the class of beliefs to which this strategy most naturally applies are beliefs that concern manifestly observable objects and qualities, such as hands and colors. Yet moral responsibility is not manifest to observation in this way. We do not see, except by a generous metaphor, that a certain person is blameworthy. Rather, this seems to be something we infer, on the basis of a host of considerations.

This reason is not decisive. We might argue that the denial of humility need not be limited in this way: it might be extended beyond the manifestly observable, perhaps so far as to include moral responsibility within its ken. Alternatively, we might say that moral responsibility is in some sense manifest, at least enough so that the epistemological considerations that tell against the external world skeptic apply also to the present case.[10]

---

[10] This thought might be licensed, for instance, by the 'phenomenal conservatism' of Michael Huemer, according to which: 'If it seems to S as if P, then S thereby has at least prima facie justification for believing that P' [2001: 99]. This is similar to Pryor's view but arguably broader in the claims that it might justify, depending on how exactly 'seeming' is construed. I consider another way of developing this thought in the next section.

The second reason is one already raised. This is the epistemic priority of excuse, which again is:

> If K is an excusing condition for some particular action A of some particular agent R, then one cannot know on partly non-testimonial grounds that R is morally responsible for A unless one independently knows that K does not obtain.[11]

The thought that grounds this claim is in the first place moral, rather than epistemological. Part of our thinking about responsibility seems to involve bestowing this sort of epistemic standing on excuses. It is not itself a principle derived from epistemological considerations, and it is not the sort of consideration that can be trumped by the epistemological considerations just adduced. No matter how sympathetically or broadly we wish to extend reasoning of the sort endorsed by Pryor, the epistemic priority of excuse gives us independent consideration against extending it as far as the proponent of the argument from moral responsibility needs it to extend, a consideration that is special to the domain of moral responsibility. This then is another reason for concluding that the grounds for denying humility in the case of external objects are not plausible grounds for denying it in the case of moral responsibility.

The alleged analogy then fails. The argument of the previous section does not prove too much, in the sense that it does not require us to endorse the reasoning of the external world skeptic. For there are principled reasons for rejecting the reasoning of the external world skeptic, reasons that do not carry over to the argument from moral responsibility. We have thus answered the challenge of explaining why the libertarian's reasoning is suspect in a way that the denial of external world skepticism is not.

---

[11] Alternatively, if it is justification rather than knowledge that is at issue (as it is, for instance, in Pryor's discussion), then the relevant claim is the modest epistemic priority of excuse.

4. Suggestions for progress

Given certain premises, the purported epistemic position of the libertarian is an impossible one. These premises are not indubitable. Indeed, we have just seen a couple of ways in which analogous principles in other domains, at least, might be cast in doubt. These considerations do not seem to apply, for reasons also just given, to the case of moral responsibility. There yet remain some epistemic approaches that the libertarian might adopt in the face of these arguments. I close by sketching, very broadly, some ideas about the forms these might take.

At least three approaches seem viable. The first two of them correspond to the two approaches developed against the external world skeptic in the previous section, namely denying particularity in light of considerations about explanation, and denying humility in light of considerations about the epistemology of perceptual experience. The third and last of them involves denying authority. Let us take these in turn.

One way to resist external world skepticism was to find an explanatory role for external objects, and to claim to know of them not particularly but rather by way of an inference to the best explanation. It may yet be that the libertarian can find a explanatory role for facts about moral responsibility. Some philosophers, such as Uri Leibowitz [2011], have stressed the existence and significance of distinctively moral explanations. Perhaps considerations about moral explanation more broadly could ground some way of understanding moral responsibility itself as having an explanatory role. This has not, to my knowledge, been done, but it is a path open for the libertarian.

The other way to resist external world skepticism was to find a perceptual role for external objects, and to claim that our knowledge of them was perceptual and hence

immune to the condition of humility. It may yet be that the libertarian can find a perceptual place for facts about moral responsibility. One way of proceeding that seems particularly promising is as follows. Since P.F. Strawson [1962] at least, attention has been focused on the connection between moral responsibility and emotions such as resentment. Perhaps the emotions offer deliverances not unlike the deliverances of external sense perception, and they too might have an epistemic role to play. If that is correct, then resentment and its cousins may immediately present us with facts about moral responsibility, so that the libertarian could take an approach to these epistemic issues very much like the approach to external world skepticism described in the foregoing. This too has not, to my knowledge, been done, but it is another path open for the libertarian.

The third path for the libertarian is simply to deny a principle that was not questioned in the case against external world skepticism, but which is perhaps the weakest of the premises in the case for the impossibility result. This was authority, the principle that someone can know that some agent is morally responsible only if she can know that some agent is morally responsible on partly non-testimonial grounds. The prospects for a thoroughly testimonial account of the epistemology of moral responsibility seem worth exploring. There is, however, a dilemma that confronts anyone who wishes to deny authority. Let us say that, contrary to authority, someone might have knowledge of moral responsibility without having any knowledge of moral responsibility that is partly non-testimonial. Her knowledge of moral responsibility will then be entirely based on testimony. The dilemma arises as follows. Consider the person from whom she has gotten this testimony: either that person is her epistemic peer, or is not. If she is, then the arguments of the foregoing apply also to her, and the concern about authority applies

equally to her. On this horn of the dilemma, the challenge from authority is simply, as it were, deferred. If the provider of testimony is not an epistemic peer, then we are on the other horn of the dilemma. On this horn of the dilemma, the libertarian owes us an explanation of how precisely it is that someone is able to transcend the epistemic limitations that grounded the impossibility result, so as to provide independently valuable testimony on the question of moral responsibility.

There are difficulties for all of these responses. Perhaps these difficulties can be resolved only by further development of these approaches. And no doubt there are other ways for the libertarian to respond. Whether or not any of these responses is viable, and indeed whether or not libertarianism is in the end a position that can command reasonable assent, it would be enlightening to know which of these forms, if any, the epistemology of moral responsibility ought to take.[12,13]

References

Fischer, John Martin and Ravizza, Mark 1998. *Responsibility and Control: A Theory of Moral Responsibility,* Cambridge: Cambridge UP.

Harman, Gilbert 1977. *The Nature of Morality*, Oxford: Oxford UP.

Hills, Alison 2009. Moral Testimony and Moral Epistemology, *Ethics* 120: 94-127.

[12] The development of such an epistemology might also be relevant to addressing other, rather different, challenges to claims about moral responsibility, such as that of Rosen 2004.

[13] Earlier versions of this essay were presented at Charles Sturt University and at a meeting of the Australasian Association of Philosophy. I would like to thank the audiences on both occasions for very helpful discussions. Thanks also to anonymous referees for this journal, one of them especially for his or her extensive and thorough comments.

Huemer, Michael 2001. *Skepticism and the Veil of Perception*. Lanham, Maryland: Rowman & Littlefield.

Leibowitz, Uri 2011. Scientific Explanation and Moral Explanation. *Noûs* 45: 472-503.

Kane, Robert 1996. *The Significance of Free Will*. Oxford: Oxford UP.

Moore, G.E. 1939. Proof of an External World, in *G.E. Moore: Selected Writings* (1993), ed. Thomas Baldwin, London: Routledge.

Pryor, James 2000. The Skeptic and the Dogmatist, *Noûs* 34: 517-549.

Rosen, Gideon 2004. Skepticism about Moral Responsibility, *Philosophical Perspectives* 18: 295-313.

Sehon, Scott 2012. Epistemic Issues in the Free Will Debate: How Can We Know When We Are Free? *Philosophical Studies* (Published online: 30 October 2012).

Strawson, P.F. 1962. Freedom and Resentment. *Proceedings of the British Academy* 48: 1-25.

van Inwagen, Peter. 1983. *An Essay on Free Will*, Oxford: Clarendon Press.

Wallace, R. Jay. 1994. *Responsibility and the Moral* Sentiments, Cambridge, Mass.: Harvard UP.