

The Metaphysical Irrelevance of the Compatibilism Debate (and, More Generally, of Conceptual Analysis)

Mark Balaguer
California State University, Los Angeles

Abstract

It is argued here that the question of whether compatibilism is true is irrelevant to metaphysical questions about the nature of human decision-making processes—for example, the question of whether or not humans have free will—except in a very trivial and metaphysically uninteresting way. In addition, it is argued that two other questions—namely, the conceptual-analysis question of what free will *is* and the question that asks which kinds of freedom are required for moral responsibility—are also essentially irrelevant to metaphysical questions about the nature of human beings.

[Compatibilism] is a wretched subterfuge ... and ... a petty word-jugglery.

—Immanuel Kant (1788, 95–96)

[Compatibilism] is a quagmire of evasion under which the real issue of fact has been entirely smothered.... No matter what the soft determinist mean by ["freedom"],... there *is* a problem, an issue of fact and not of words.

—William James (1884, 149)

*Mark Balaguer is Professor of Philosophy at California State University, Los Angeles. His areas of interest are metaphysics, philosophy of mathematics, philosophy of language, free will, and metaethics. He is the author of *Platonism and Anti-Platonism in Mathematics* (Oxford University Press, 1998), as well as numerous journal articles.*

1. Introduction

The recent literature on the problem of free will has been dominated by the debate about whether compatibilism is true. I think this is unfortunate because I think the question of whether compatibilism is true is essentially irrelevant to metaphysical questions about human free will—that is, about the existence or nature of the freedom inherent in human decision-making processes. This view of the compatibilism debate is *metaphysically* similar to the views of Kant and James, but I should say here that I would not go along with the dismissiveness of their remarks. I think the question of whether compatibilism is true is interesting in its own right, and as will become clear below, I think it may be worth caring about for reasons that go beyond our purely theoretical interests. Moreover, I think there has been a lot of good philosophical work on this question that is both interesting and important. But I also think that upon reflection, it turns out that the compatibilism question is independent of metaphysical questions about the nature of human decision-making processes, except in a trivial way. In particular, I will argue that

The compatibilism question: Is free will compatible with determinism?¹

is essentially irrelevant to

The do-we-have-free-will question: Do human beings have free will?

I will do this by arguing for a few different theses. To begin with, I will argue that the two questions listed here—that is, the compatibilism question and the do-we-have-free-will question—reduce to two other questions that can be thought of as more fundamental, namely,

The what-is-free-will question: What is free will? (We can take this as being equivalent to the question ‘What is the correct analysis of the notion of free will?’ and also to the question ‘What is the correct definition of the term “free will”?’ But we cannot assume without argument that these questions are solely about folk meaning, or ordinary-language usage and intentions; I will discuss this issue below.)

and

The which-kinds-of-freedom-do-we-have question: Which kinds of freedom do humans have? That is, do they have libertarian freedom (or for short, *L-freedom*²); and do they have Humean

freedom?; and do they have Frankfurtian freedom?; and so on.³ (Actually, to be more precise, we can formulate this question as asking which kinds of “freedom” humans have, since some or all of the kinds of “freedom” we’re asking about here might fail to *be* free will, according to the correct answer to the what-is-free-will question.)

If we could answer these two questions, then (I will argue) we could also answer the compatibilism question and the do-we-have-free-will question. In particular, I will argue that (i) the compatibilism question reduces to the what-is-free-will question, and (ii) the do-we-have-free-will question collapses into the what-is-free-will question *and* the which-kinds-of-freedom-do-we-have question.

I will argue point (i) below, in section 5, but point (ii) is more or less obvious and can be motivated right here in just a few words: if we could answer the what-is-free-will question and the which-kinds-of-freedom-do-we-have question, then we could also answer the do-we-have-free-will question, because we would know what free will is and whether we have it. Of course, we would also know more than this—in particular, we would know whether we possess various kinds of pseudo-freedom—but this doesn’t undermine the claim that the do-we-have-free-will question is subsumed by the what-is-free-will question and the which-kinds-of-freedom-do-we-have question.

The main thing I want to argue, however, is this:

(iii) While the what-is-free-will question is clearly relevant to the do-we-have-free-will question in a certain way, it is not relevant to that question in any nontrivial or metaphysically interesting way; indeed, the what-is-free-will question is not relevant to *any* substantive question about the nature of human decision-making processes, except in a trivial way.

I will argue this point in sections 3 and 4. And if we combine thesis (iii) with thesis (i)—which, again, I’ll argue in section 5—we obtain the result that the compatibilism question is also essentially irrelevant to substantive questions about the nature of human decision-making processes, for instance, the do-we-have-free-will question. Moreover, as we’ll see, my argument points to a more general result, namely, that conceptual analysis is essentially irrelevant to metaphysics.

If my arguments are correct, then what we might think of as the *semantic* component of the problem of free will (namely, the what-is-free-will question) is essentially irrelevant to the *metaphysical* component of the problem (e.g., the do-we-have-free-will question and the which-kinds-of-freedom-do-we-have question). But this isn’t all there is to the problem of free will. There is also a *moral* component to the problem, for we also

want to know which kinds of freedom are required for moral responsibility. But I will argue in section 4 that however important this question is to moral issues, it is—like the compatibilism question and the what-is-free-will question—essentially irrelevant to metaphysical questions about the nature of human decision-making processes, for example, the do-we-have-free-will question.

2. What Determines Whether an Answer to the What-Is-Free-Will Question Is Correct?

Before I construct my main argument (in sections 3–5), I need to discuss an important question: What determines whether a given answer to the what-is-free-will question is correct? There are numerous views one might adopt here. Here's one very simple view:

The OL View: An answer to the what-is-free-will question is correct iff it captures the ordinary-language meaning(s) of the expression 'free will'—that is, what this expression means among ordinary folk. (Of course, it doesn't follow from this that ordinary usage of 'free will' picks out a unique, well-defined concept; proponents of the OL View can allow that there might be vaguenesses, inconsistencies, ambiguities, and so on built into our usage here; but they would say that if the ordinary term 'free will' is indeed vague or ambiguous, then a complete and correct answer to the what-is-free-will question would tell us about this.)

This view is at least initially plausible. One point that counts in its favor is that (a) it makes the what-is-free-will question *factual* (it is undeniable, I think, that there are facts about how ordinary folk use expressions like 'free will', 'can', 'could have done otherwise', and so on); and (b) it's not obvious that any other plausible view does this. A second point that counts in favor of the OL View is that it fits very well with the methodology that philosophers actually use when they try to answer questions like 'What is free will?', 'What is knowledge?', and so on. One of the main things philosophers do here is use ordinary-language intuitions (about the applicability of our concepts in real and imagined scenarios) to confirm and falsify theories; for instance, if a given theory of free will—that is, a given answer to the what-is-free-will question—flies in the face of our intuitions about when the concept of free will does and doesn't apply, this is seen as falsifying evidence. But it seems to me that this methodology would make little sense if we weren't at least partially engaged in trying to uncover ordinary-language meaning. For while it's plausible to suppose that our intuitions

reliably track facts about ordinary-language meaning, it's not very plausible to suppose that they reliably track other kinds of facts.

By the way, I should say that the OL View is neutral on the question of whether, in trying to answer the what-is-free-will question, we should be doing real empirical studies on folk intuitions. According to one traditional view, we don't need to do this, because we are ourselves native speakers of ordinary language, and so we can use our own intuitions. But there is another view that's growing in prominence (see, e.g., Nichols 2006 and Nahmias, et al. 2005) that holds that our answers to questions like the what-is-free-will question should be based on real empirical data about the intuitions of ordinary folk, in particular, nonphilosophers. The OL View is consistent with both of these methodological views, and the issue here will not be relevant to what I will be arguing.

In any event, it seems to me that the OL View is at least initially plausible, and it might even be the right view. But on the other hand, it might not be, for there are other views one might endorse here. For instance, one might maintain that while it is important to uncover the facts about what ordinary folk mean by the term 'free will', this isn't all there is to arriving at a fully acceptable answer to the what-is-free-will question. One might think that other issues need to be considered as well, for example, issues having to do with the coherence of ordinary-language conceptions of freedom, or with the kinds of freedom that are required for moral responsibility, or with the kinds of freedom that are actually at work when human beings have the experience of acting and choosing freely.

Here's what I'm going to do: in section 3, I will assume that the OL View is correct, and based on that assumption, I will argue that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question. Then in section 4, I will argue that even if we drop the assumption that the OL View is correct, we still get the same conclusion about the metaphysical irrelevance of the what-is-free-will question. Indeed, I will argue that we get this result no matter what view we adopt of the correctness conditions of answers to the what-is-free-will question—that is, no matter what we say about what determines whether an answer to the what-is-free-will question is correct.

3. Why the What-Is-Free-Will Question Is Irrelevant to the Do-We-Have-Free-Will Question, Assuming the OL View Is Correct

In this section, I will assume that the OL View is correct and argue for the following thesis:

(iii) While the what-is-free-will question is clearly relevant to the do-we-have-free-will question in a certain way, it is not relevant to that question in any nontrivial or metaphysically interesting way; indeed, the what-is-free-will question is essentially irrelevant to *all* substantive questions about the nature of human decision-making processes.

Before giving my argument, I want to say what I mean by 'metaphysically interesting'. In the present context, I'm using 'metaphysical' to mean something like *about the world*; thus, since the topic here is human decision-making processes and human freedom, 'metaphysical' is being used to mean something like *about human beings or human decision-making processes*. Thus, if I say that a question is not relevant to the do-we-have-free-will question in any metaphysically interesting way, what I mean is essentially that it's not relevant to that question in a way that's substantively relevant to the goal of discovering the nature of human beings and, in particular, human decision-making processes. And note that I'm not claiming here that this is what anyone else means by 'metaphysically interesting'. I'm just specifying what *I* mean. Thus, if someone were to respond that, on their view, conceptual analysis is a *part* of metaphysics and, hence, that it's obviously metaphysically interesting, my response would be to give them the expression 'metaphysically interesting'. I would just rephrase my thesis in terms of relevance to the goal of discovering the nature of human beings and human decision-making processes.

I want to begin my argument for (iii) by saying why someone might want to reject it. Thus, consider the following counter-argument or objection to (iii):

Until we determine what free will *is*, we can't determine whether humans have free will because we won't even know what we're talking about, or looking for. Thus, the what-is-free-will question is obviously relevant to the do-we-have-free-will question, and moreover, philosophers who are engaged in trying to answer the what-is-free-will question are *doing metaphysics* because they're doing something that's centrally important to the task of answering the do-we-have-free-will question, which is clearly a metaphysical question about the nature of human beings.

I think the central claims in this objection are essentially right, but I also think they overlook an important point, and it's for this reason that thesis (iii) contains a proviso about metaphysical interestingness. What I want to argue here is that when we take this proviso seriously, we see that thesis (iii) is importantly correct. More precisely, I want to argue this point on the assumption that the OL View is correct—that is, on the assumption

that an answer to the what-is-free-will question is correct iff it captures the ordinary-language meaning(s) of 'free will'.

Let me begin here by changing the example and showing how awkward the line of thought in the above objection would be if we employed it in other settings. To this end, suppose that Carstairs and Caruthers are linguists (or sociologists or psychologists or whatever) who specialize in trying to figure out the precise meanings of scientific terms, as they're actually used by working scientists (and perhaps ordinary folk); and suppose in particular that Carstairs and Caruthers have competing theories of the ordinary-discourse meaning of the word 'planet'. Finally, suppose that the dispute between Carstairs and Caruthers has implications for the truth values of various astronomical sentences. For example, suppose that astronomers have discovered an object, call it Wilma, that's orbiting the sun beyond Pluto and that counts as a planet if Carstairs's theory of the meaning of 'planet' is right but doesn't count as a planet if Caruthers's theory is right, so that the truth value of the sentence 'There is a tenth planet in our solar system' depends on whether Carstairs or Caruthers is right (of course, if you don't think Pluto counts as a planet, then the question would be whether there's a *ninth* planet in the solar system, but let's ignore this complication). Now, if we take the style of thinking inherent in the above objection and apply it in the present case, we seem to obtain the result that when Carstairs and Caruthers are debating what the ordinary meaning of 'planet' is, they're doing *astronomy*. But, of course, they're not doing astronomy; they're doing empirical semantics. Their investigation isn't relevant in any nontrivial way to an inquiry into the nature of the solar system. If astronomers know that Wilma is there, and if they know how big Wilma is, and what it's made of, and what its orbital path is like, and so on and so forth, then their work is done. It would be a bizarre, misleading representation of the situation to claim that Carstairs and Caruthers were disputing an open astronomical question, that is, a question about the nature of the solar system, and that they were trying to settle this question by studying the ordinary usage of astronomers.

Similar remarks, it seems, can be made about free will. Philosophers involved in trying to answer the what-is-free-will question are not engaged in a genuine inquiry into the nature of human-decision making processes, and indeed, their investigations are not relevant to such inquiries in any nontrivial way. Like Carstairs and Caruthers, what these philosophers are doing is empirical semantics; they're engaged in an investigation of ordinary-language usage and intentions.

Let's change the example now. Let's suppose it's an open astronomical question whether Wilma exists. For instance, we can suppose that some astronomer has discovered some pertur-

bations in Pluto's orbital path and hypothesized Wilma's existence to explain these perturbations but that, as of yet, no one has actually found Wilma. And let's suppose that a dispute has arisen over this issue and that people characterize this dispute by saying that there's a controversy among astronomers over the following:

The tenth-planet question: Is there a tenth planet in our solar system (or a ninth planet, if you think that Pluto isn't a planet)?

Now suppose that Caruthers announces that he has discovered the answer to the tenth-planet question; he argues that Wilma may or may not exist—he has no idea about this—but that it doesn't matter because even if Wilma does exist, it is not a planet, and so the answer to the tenth-planet question is 'No'. Finally, suppose that Carstairs argues that Caruthers is mistaken; he says that if Wilma exists then it *is* a planet, and so, he says, we have to determine whether Wilma exists in order to answer the tenth-planet question. Now, what I want to ask is this: Is the Carstairs–Caruthers debate relevant to the tenth-planet question? Well, there is obviously a sort of relevance here because Caruthers's semantic thesis (together with various theses about the solar system that we're assuming are endorsed by all the parties to the dispute) entails that the answer to the tenth-planet question is 'No'. But from an *astronomical* point of view—that is, from the point of view of the goal of discovering the nature of the solar system—this is clearly a trivial sort of relevance. When we ask the tenth-planet question, we are presumably trying to learn something about the nature of the solar system; but the Carstairs–Caruthers debate is relevant to that goal in at most a trivial way, for how the word 'planet' happens to be used in ordinary discourse doesn't tell us anything important or nontrivial about the nature of the solar system. Therefore, it seems fair to say that while the Carstairs–Caruthers debate is obviously relevant to the tenth-planet question in a certain sort of way, it is not relevant to that question—or indeed to any question about the nature of the solar system—in any nontrivial or astronomically interesting way. (This, of course, is not to say that the Carstairs–Caruthers debate is itself trivial or uninteresting; it's just to say that it's not an astronomical debate, i.e., a debate about the nature of the solar system, except in a very trivial way.)

Likewise, if we ask whether the what-is-free-will question is relevant to the do-we-have-free-will question, we can say that there is obviously a sort of relevance here, because, for example, Hume's answer to the what-is-free-will question (together with other theses that just about all of us accept) entails that the answer to the do-we-have-free-will question is 'Yes'. But from a metaphysical point of view—that is, the point of view of the

goal of discovering the nature of human decision-making processes—this is clearly a trivial sort of relevance. When we ask the do-we-have-free-will question, we are presumably trying to learn something about the nature of human decision-making processes; but the what-is-free-will question is relevant to that goal in at most a trivial way, for how the expression 'free will' happens to be used in ordinary discourse doesn't tell us anything important or nontrivial about the nature of human decision-making processes. Therefore, once again, it seems fair to say that while the what-is-free-will question is obviously relevant to the do-we-have-free-will question in a certain way, it is not relevant to that question—or indeed to any question about the nature of human decision making—in any nontrivial way. And again, the point here is not that the what-is-free-will question is itself trivial or uninteresting; the point is that it is not a question about human beings, or human decision making, except in a trivial way; it is rather about the semantics of a certain expression.

There's another parallel between the planet case and the free will case that's worth commenting on. If what I'm arguing is correct, then the metaphysically interesting issue behind the do-we-have-free-will question is captured by the which-kinds-of-freedom-do-we-have question, which is entirely independent of the what-is-free-will question (indeed, this question was purposely formulated in a way that would make it independent of the ordinary-language meaning of 'free will'). If we could fully answer the which-kinds-of-freedom-do-we-have question, then we would have all the relevant facts about the nature of human decision-making processes that we might need in order to answer the do-we-have-free-will question; we still might not know the answer to the do-we-have-free-will question, but that would just be a function of our not knowing the ordinary-language meaning of the term 'free will'; it would not signify any substantive ignorance about the nature of human beings. Likewise, it seems that the astronomically interesting issue behind the tenth-planet question is captured by a question that's entirely independent of the Carstairs–Caruthers debate, namely, the question 'Is there any such thing as Wilma?' If we could answer this question, we would know all the relevant facts about the nature of the solar system; we still might not know the answer to the tenth-planet question, but that would just be a function of our not knowing what 'planet' means; it would not signify any substantive ignorance about the nature of the solar system.⁴

The point here can be generalized. Whenever you're trying to discover something about the nature of the world, you can always proceed straight to the point at hand, without having to determine the meaning of some folk expression, by simply introducing some theoretical terms and defining them by stipulation. Thus, for example, if you just want to know what the

solar system is like, you can forget about folk terms like ‘planet’ and introduce some new terms with clearly defined meanings. And if you just want to know what human decision-making processes are like, you can simply use terms of art like ‘Humean freedom’ and ‘L-freedom’ and so on and proceed straight to the point at hand, trying to determine which of the various kinds of freedom (or “freedom”) human beings actually possess without first determining the ordinary-language meaning of the folk term ‘free will’. And if you’re in a situation where you already know all the relevant metaphysical facts but don’t know what some folk term means, then you can describe the metaphysical facts using technical terms with stipulated definitions, and so your lack of knowledge of the meaning of the folk term shouldn’t be treated as a genuine ignorance of (nonsemantic) metaphysical facts. (Another way to appreciate the generality of the issue here is to note that behind *every* question about the nature of the world, there are semantic questions about the meanings of the words in the given about-the-world question, and these semantic questions are, in some sense, logically prior to the question about the nature of the world; but insofar as one is interested in discovering the nature of the *world* when one is asking the about-the-world question, and not in discovering the meanings of the words in that question, there is no need to trouble oneself with the semantic questions before addressing the issue about the nature of the world.)

Before going on, I want to draw a distinction between two different aspects of conceptual analysis and acknowledge that one of them is in fact relevant to metaphysical questions about human beings (but in a way that doesn’t undermine my position here). The distinction I have in mind is between (a) the *articulation* of concepts, or conceptual analyses, and (b) arguments for and against the *correctness* of the various analyses that have been offered. The former is definitely relevant to metaphysical questions about the nature of human decision-making processes: every time someone comes up with a new analysis of free will, it generates a new subquestion of the which-kinds-of-freedom-do-we-have question, and so this can lead to the discovery of new and interesting facts about humans. What I’m claiming is *not* relevant (in any nontrivial way) to metaphysical questions about human beings is the question of which of the many analyses of free will that people have articulated are *correct*. And this, of course, is just what the what-is-free-will question asks.

4. Why the What-Is-Free-Will Question Is Irrelevant to the Do-We-Have-Free-Will Question, Even If the OL View *Isn’t* Correct

I just argued that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question. In arguing this

point, I assumed the OL View; that is, I assumed that an answer to the what-is-free-will question is correct iff it captures the ordinary-language meaning(s) of ‘free will’. I now want to argue that even if we reject the OL View, we still get the result that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question.

If the OL View is wrong, then we can’t answer the what-is-free-will question by looking only at facts about ordinary-language meaning; we need to look at other kinds of facts as well (or instead, as the case may be). Indeed, one might think that in order to adequately answer the what-is-free-will question, we need to take a variety of different issues into account. One way to motivate this stance would be to endorse something like the following view:

The Hybrid View: It’s true that *part* of what we’re doing in trying to answer the what-is-free-will question is trying to figure out what ordinary folk mean by ‘free will’. But that’s not all we’re doing. We’re also trying to improve upon folk usage, or supplement it, or some such thing. In doing this, we might need to consider a number of different issues. For instance, we might want to eliminate some incoherence or imprecision from the ordinary folk notion of free will. Or, second, we might want to take into account the issue of what’s *worth wanting*; in other words, in order to determine what free will is, we might need to figure out which kinds of freedom are *valuable*; in particular, we might need to figure out which kinds of freedom are required for moral responsibility, or autonomy, or dignity, or other things we might value. Or, third, we might want to take into account the kinds of freedom that humans actually have, for one might think that what free will *is* is at least partially determined by the kinds of freedom that are actually at work when people have the experience of acting and choosing freely.

Regardless of whether the Hybrid View, as its stated here, is the best alternative to the OL View, it seems that the sorts of considerations mentioned in the Hybrid View—most notably, those having to do with coherence, moral responsibility, and the nature of actual human freedom—are the most plausible candidates for what might be relevant to the what-is-free-will question, if that question isn’t simply about ordinary-language meaning. Thus, it seems to me that if we reject the OL View, we’re going to wind up saying that questions like the following are relevant to the what-is-free-will question:

The which-kinds-of-freedom-do-we-have question (see section 1 for a formulation).

The coherence question: Which kinds of freedom (or “freedom”) are coherent, or conceptually possible?

The moral responsibility question: Which kinds of freedom (or “freedom”) are required for moral responsibility?

Now, I think there are potential problems with the view that questions like these are relevant to the what-is-free-will question—for instance, one might argue that the OL View is the best view after all—but I don’t want to pursue this issue here. Instead, I want to argue that even if we assume that questions like these *are* relevant to the what-is-free-will question, we still get the result that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question (and, indeed, to other questions about the nature of human decision-making processes). I will argue this point in connection with the three questions listed above in sections 4.1–4.3. Then in section 4.4, I will construct a more general argument, one that motivates the idea that no matter what we say about the kinds of facts and questions that might be relevant to the what-is-free-will question, we still get the result that that question is essentially irrelevant to the do-we-have-free-will question.

4.1 The Which-Kinds-of-Freedom-Do-We-Have Question

There are a few different reasons one might give for thinking that the which-kinds-of-freedom-do-we-have question is relevant to the what-is-free-will question. For instance, one might think that our predicates should be interpreted in ways that enable them to “carve nature at its joints,” and given this, one might argue that we ought to take ‘free will’ to refer to a kind of freedom that humans actually have—that is, a kind that’s actually present when people have the experience of acting and choosing freely. Or alternatively, one might argue that ‘free choice’ is a kind term, or a paradigm-case term, and one might infer from this that ‘free will’ refers to the kind of freedom that’s inherent in ordinary human choices, whatever that turns out to be. I have serious doubts about both of these views, and more generally, I have doubts about the claim that the which-kinds-of-freedom-do-we-have question is relevant to the what-is-free-will question. But again, I don’t want to pursue this here.⁵ Instead, I want to argue that even if the which-kinds-of-freedom-do-we-have question *is* relevant to the what-is-free-will question, this doesn’t undermine my claim that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question. The argument for this, as we’ll presently see, is very simple.

The view I’ve been advancing here is that (a) the do-we-have-free-will question decomposes into the which-kinds-of-

freedom-do-we-have question and the what-is-free-will question, and (b) the which-kinds-of-freedom-do-we-have question contains everything that’s metaphysically interesting here, so that the what-is-free-will question doesn’t add anything of metaphysical interest. Now, if it turns out that the which-kinds-of-freedom-do-we-have question is relevant to the what-is-free-will question, then the latter question is obviously a metaphysically interesting question. But it’s not metaphysically interesting in any *new* way. For the metaphysically interesting facts that would be relevant to the what-is-free-will question in this scenario are the very same facts that are relevant to the which-kinds-of-freedom-do-we-have question, and so they’re the same facts that are *already* relevant to the do-we-have-free-will question.

In short, the point is this: even if we assume that the which-kinds-of-freedom-do-we-have question is relevant to the what-is-free-will question, we still get the result that if we could answer the which-kinds-of-freedom-do-we-have question, we would have all the relevant metaphysically interesting information we would need in order to answer the do-we-have-free-will question—and so we still get the result that the what-is-free-will question adds nothing here that’s metaphysically interesting. In other words, we still get the result that if we could already answer the which-kinds-of-freedom-do-we-have question, then by moving on and answering the what-is-free-will question, we would not be learning anything new about the nature of human beings or human decision-making processes. And so our situation here isn’t changed at all by assuming that the which-kinds-of-freedom-do-we-have question is relevant to the what-is-free-will question.

Another way to put the point here is in terms of direction of relevance: if the which-kinds-of-freedom-do-we-have question were relevant to the what-is-free-will question, this wouldn’t give us the result that the what-is-free-will question was relevant to metaphysical questions about human decision making; rather, it would give us the opposite result—that is, that metaphysical questions about human decision making were relevant to the what-is-free-will question. In this scenario, we could independently discover facts about humans that turned out to be relevant to determining what free will *is*, but we could not independently discover what free will *is* in a way that would make this relevant to figuring out which sorts of freedom human beings actually possess, for in this scenario, in order to figure out what free will *is*, we would first have to figure out which sorts of freedom humans actually have.

4.2 The Coherence Question

A very similar argument can be used in connection with the coherence question. Now, I want to admit that the coherence

question is itself a metaphysically relevant question. After all, if a kind of freedom is incoherent, or conceptually impossible, then it follows that humans do not possess that kind of freedom. Thus, the coherence question is directly relevant to the which-kinds-of-freedom-do-we-have question, and so it is a metaphysically relevant question. But the coherence question is not metaphysically relevant in any way that undermines my position here. Recall my thesis: (a) if we could answer the which-kinds-of-freedom-do-we-have question, then we would have all the relevant metaphysically interesting information we would need in order to answer the do-we-have-free-will question, and (b) by going on and answering the what-is-free-will question, we wouldn't be learning anything that was both new and metaphysically interesting—that is, we wouldn't be learning anything new about the nature of human decision-making processes. But this will be true even if we assume that the coherence question is relevant to the what-is-free-will question. For, in short, the metaphysically relevant features of the coherence question are *already* relevant to the which-kinds-of-freedom-do-we-have question (and, hence, to the do-we-have-free-will question). If we could already answer the which-kinds-of-freedom-do-we-have question, and if some specific kind of freedom F were conceptually impossible, or incoherent, then we would *already know* that we didn't possess F. Thus, even if we could use this information about F in answering the what-is-free-will question—for example, even if we could use it to determine that free will isn't F—in doing this, we wouldn't be learning anything new about human beings or human decision making. For, again, we would already know all the relevant metaphysically interesting facts. Thus, again, even if we assume that the coherence question is relevant to the what-is-free-will question, we still get the result that the what-is-free-will question adds nothing of metaphysical interest to the which-kinds-of-freedom-do-we-have question.

4.3 The Moral Responsibility Question

Once again, I have serious doubts about the thesis that the moral responsibility question is substantively relevant to the what-is-free-will question, but again, I don't want to pursue this here. Instead, I want to argue that even if we assume that the moral responsibility question is relevant to the what-is-free-will question, it doesn't matter (in the present context) because the moral responsibility question is itself essentially irrelevant to metaphysical questions about the nature of human decision-making processes.

One might argue that the moral responsibility question is relevant to metaphysical questions about humans by saying something like this:

Suppose we knew which kinds of freedom (or “freedom”) humans possessed but didn't know which of them were required for moral responsibility. Then we wouldn't know whether humans were morally responsible for their actions, and so by answering the moral responsibility question, we could discover a fact *about humans*—namely, whether they were morally responsible for their actions.

I have two responses to this argument. First, even if the moral responsibility question is relevant to a question about humans (namely, the question of whether humans are morally responsible for their actions), in the present context, this isn't the right kind of metaphysical relevance. In order to undermine my position, we would need the result that the moral responsibility question is relevant to questions about the nature of human decision-making processes, and it's hard to believe that this is true. If the moral responsibility question is relevant to the what-is-free-will question, then we might get the result that the what-is-free-will question is relevant to questions about the moral status of humans. But we don't get the result that it's relevant to questions about the metaphysical nature of human decision-making processes; for, again, by answering the moral responsibility question and, hence, the what-is-free-will question, we wouldn't be learning anything new about the nature of our decision-making processes, in particular, about the kinds of freedom (or “freedom”) that we actually have.

The second response to the above argument is that, perhaps surprisingly, the moral responsibility question is in fact not relevant (in any metaphysically interesting way) to the question of whether human beings are morally responsible for their actions. To argue this point, let me assume, for the sake of simplifying things, that (a) aside from the issue of the kinds of freedom that are required for moral responsibility, we know roughly what moral responsibility is, and (b) we also know that humans do satisfy all the requirements for moral responsibility except possibly for the freedom requirement. (Of course, we don't really know these things, but I won't be begging any questions by assuming them here. Indeed, it's easy to see that if humans fail to satisfy some other requirement for moral responsibility—that is, some nonfreedom requirement—then the above objection to my view completely falls apart because we would already know, independently of the free will issue, that humans aren't responsible for their actions.) In any event, given these assumptions, let's suppose that we already knew the answer to the which-kinds-of-freedom-do-we-have question but didn't know the answer to the moral responsibility question, and hence, didn't know whether we had free will or moral responsibility, because we didn't know what free will or moral responsibility were. Then we could define a series of kinds of moral

responsibility (or “moral responsibility,” as the case may be) corresponding to the different kinds of freedom (or “freedom”). For instance, if we had named the various kinds of freedom F_1 , F_2 , and so on, then R_1 would be a kind of moral responsibility that required F_1 (and no other kind of freedom); R_2 would be a kind of moral responsibility that required F_2 (and no other kind of freedom); $R_{7,12}$ would be a kind of moral responsibility that required F_7 and F_{12} (and no other kind of freedom); and so on. It follows from everything we’re assuming here that in this scenario, we would already know which of these kinds of responsibility we possessed and which we didn’t. So we would already know all the responsibility facts about humans. We might not know the truth value of the sentence ‘Humans are morally responsible for their actions’, but that wouldn’t be a function of any lack of knowledge *about humans*. It would simply be because we didn’t know what moral responsibility *was*, or what ‘moral responsibility’ meant, or some such thing. Thus, it seems to me that the moral responsibility question isn’t relevant, in any metaphysically interesting way, to the question of whether humans are morally responsible for their actions—or, indeed, to any question about humans.

To give a concrete example here, my claim is this: suppose that (a) we already knew that humans were Hume-responsible but not libertarian-responsible, but (b) we didn’t know whether moral responsibility was Hume-responsibility or libertarian-responsibility (although we can assume that we had figured out that it was one of the two and not some third kind of responsibility); then we wouldn’t be lacking any knowledge here about humans, in any interesting sense; we would simply be lacking knowledge of what moral responsibility is.

What sort of knowledge is this that we would be lacking? Well, one view is that it’s essentially just knowledge of ordinary-language meaning. On this view, to say that moral responsibility is, for example, libertarian-responsibility is to make a claim about the ordinary concept of moral responsibility, that is, about the way that ordinary folk use the term ‘moral responsibility’. On this view, for all the reasons given in section 3, the moral responsibility question is not relevant in any metaphysically interesting way to any questions about humans.

What other sort of question might the moral responsibility question be, if not a question about ordinary-language meaning? Well, the only other plausible view, I think, is that it’s partially about capturing ordinary-language meaning and partially about improving on ordinary-language meaning. One way to improve on ordinary meaning would be to eliminate incoherences, but I could presumably take the same line here that I took on this issue in section 4.2. Another way to improve on the ordinary usage of ‘moral responsibility’ would be to figure out which of the various kinds of moral responsibility are *fair* (where a kind

of responsibility is fair iff it would be fair to hold someone morally responsible because they had that kind of responsibility). But it’s hard to believe that we could make any progress here by appealing to the notion of fairness because, presumably, each kind of responsibility brings with it its own kind of fairness. Now, of course, we might try to determine which kinds of responsibility are *really fair*, but it seems to me that this could only be a question about the *ordinary* notion of fairness, and so this would put us right back in the business of determining ordinary-language meaning. So, again, it seems to me that there’s no way to get the result that the moral responsibility question is relevant, in a metaphysically interesting way, to questions about the nature of human beings.

4.4 Generalizing the Argument

I now want to construct a general argument for thinking that no matter what we say about the kinds of facts that might be relevant to the what-is-free-will question—that is, the kinds of facts that might determine which answer to the what-is-free-will question is correct—we still get the result that the what-is-free-will question is essentially irrelevant to the do-we-have-free-will question (and, indeed, to all questions about the nature of human decision-making processes). There are two sorts of facts that one might think relevant to the what-is-free-will question, namely, (i) facts that are relevant to the which-kinds-of-freedom-do-we-have question, and (ii) facts that aren’t relevant to that question, for example, facts about ordinary-language meaning. But as we saw in section 4.1, type-(i) facts aren’t relevant in any *new* way to the do-we-have-free-will question, because we already have to answer the which-kinds-of-freedom-do-we-have question in order to answer the do-we-have-free-will question. And, second, it’s hard to see how type-(ii) facts could be relevant, in any metaphysically interesting way, to the do-we-have-free-will question. For the which-kinds-of-freedom-do-we-have question seems to capture *all* the metaphysically interesting facts that might be relevant to the do-we-have-free-will question. In other words, once the which-kinds-of-freedom-do-we-have question has been answered, there are simply no facts left to discover that are both relevant to the do-we-have-free-will question and metaphysically interesting in the sense of being about the nature of human beings. Thus, the conclusion I want to draw is that no matter what we say about the kinds of facts that might be relevant to the what-is-free-will question, that question is not relevant in any metaphysically interesting way to the do-we-have-free-will question. The only really metaphysically important question here is the which-kinds-of-freedom-do-we-have question.

5. Why the Compatibilism Question Reduces to the What-Is-Free-Will Question

I have now argued that the what-is-free-will question is irrelevant to the do-we-have-free-will question in all but a very trivial way. But I also want to argue that the compatibilism question is irrelevant to the do-we-have-free-will question. To establish this, I will argue in the present section that the compatibilism question reduces to the what-is-free-will question.

The argument for this is very simple. The first point to note is this: just as we found that the do-we-have-free-will question reduces to, or is subsumed by, the what-is-free-will question and the which-kinds-of-freedom-do-we-have question, so too it seems that the compatibilism question reduces to, or is subsumed by, the what-is-free-will question and what might be called the *which-kinds-of-freedom-are-compatible-with-determinism question*. The second point to be made here is this: *prima facie*, it seems that all of the subquestions of the which-kinds-of-freedom-are-compatible-with-determinism question—for example, ‘Is Humean freedom compatible with determinism?’, ‘Is L-freedom compatible with determinism?’, and so on—are trivial and obvious. In other words, it seems that all of the analyses of free will in the literature are either obviously compatible with determinism or obviously incompatible with it. For instance, L-freedom is obviously incompatible with determinism because it is by definition indeterministic (see note 2). And Humean freedom is obviously compatible with determinism because it’s essentially just the ability to act and choose in accordance with your desires, and it could be that our desires are causally determined by prior events and that these desires determine our decisions, which in turn determine our actions. And Frankfurtian freedom is obviously compatible with determinism because it’s just the ability to control, with second-order attitudes, which of your first-order desires will affect your behavior, and it could be that we are causally determined to have (and act on) second-order desires of this sort. And the same point can be made about all the different compatibilist analyses that have been put forward, for example, the proposals of people like Watson (1975), Dworkin (1988), Wolf (1990), Double (1991), Fischer (1994), Fischer and Ravizza (1998), Mele (1995), and Bok (1998), to name just a few. In sum, there are no mainstream analyses of free will—or at any rate, none that I know of—that generate substantive, nonobvious compatibility questions.⁶ Thus, it seems that the which-kinds-of-freedom-are-compatible-with-determinism question just disappears. And so the only controversial part of the compatibilism question is the what-is-free-will question—which is just to say that the compatibilism question reduces to the what-is-free-will question.

Of course, in the future, someone could come up with a new kind of freedom that generated a pure compatibility question that was controversial; this would be a purely logical question—it would be of the form ‘Is the concept of X-freedom compatible with determinism?’ But as of right now, there are no such questions that are controversial, and so we can say that, relative to the current state of the debate, the compatibilism question reduces to the what-is-free-will question.

It’s important to note that the point I’m arguing here—that the compatibilism question reduces to the what-is-free-will question—does not constitute an *objection* to any important part of the literature on the compatibilism question. Indeed, it seems to me that when we look closely at this literature, we find that it fits very nicely with my thesis. We can divide the literature on the compatibilism question into four main strands, namely, (i) the literature on the Humean conditional analysis of free will; (ii) the literature regarding recent compatibilist attempts to construct an acceptable analysis of the notion of free will; (iii) the literature on the consequence argument for incompatibilism; and (iv) the literature on the Frankfurt-case argument for compatibilism (actually, the Frankfurt-case argument is more often brought up in connection with the issue of the compatibility of determinism and moral responsibility, but all the arguments here can be reproduced in connection with the issue of the compatibility of determinism and free will—though, of course, one might doubt that the corresponding arguments always stand or fall together⁷).⁸ I think it can be argued that in all four of these strands, the really controversial issues boil down to semantic questions about the meanings of various terms (not just ‘free will’ but related expressions like ‘can’, ‘could have done otherwise’, and so on). This is more or less obvious in connection with strands (i) and (ii). It is perhaps a bit less obvious in connection with strand (iv), but not much. The issues there center around various thought experiments and intuitions, and this alone suggests that conceptual analysis is what’s at issue. The question is not whether some clearly defined concept is compatible with couldn’t-have-done-otherwise; rather, it’s whether we should interpret the term ‘moral responsibility’ (or ‘free will’) as expressing a certain sort of concept—in particular, one that’s compatible with couldn’t-have-done-otherwise. Finally, the point is probably least obvious in connection with strand (iii), that is, the consequence argument, because on the surface, the issue there doesn’t seem to be one of meaning or conceptual analysis. But upon reflection, it becomes clear that this is precisely what’s at issue. I cannot argue this point in detail, but the main idea behind the argument is as follows: (a) every version of the consequence argument is couched in terms of some crucial expression like ‘has a choice about’; and (b) if we interpret these expressions along incom-

patibilist lines, the argument is clearly sound; and (c) if we interpret them along compatibilist lines, the argument is clearly unsound; and so (d) the issue boils down to a question about how these expressions ought to be interpreted.

So, again, the arguments of this paper are not supposed to provide objections to any of the central arguments or positions in the literature on compatibilism. Moreover, I am not claiming (and I wouldn't claim) that the compatibilism question—or the what-is-free-will question or the moral responsibility question—are unimportant. Everything I've said here is perfectly consistent with the claim that these questions are extremely important. My claim is simply that they aren't relevant in any nontrivial way to the do-we-have-free-will question—or, more generally, to investigations into the nature of human decision-making processes.

Before concluding, let me make a few final points. First, if the arguments of this paper are correct, then the metaphysical component of the problem of free will essentially boils down to the which-kinds-of-freedom-do-we-have question. But notice that most of the subquestions of this question are pretty trivial. In particular, it seems more or less obvious that we do possess most of the standard compatibilist kinds of freedom, for example, Humean freedom, Frankfurtian freedom, and so on. There might be a few kinds of compatibilist freedom that are slightly controversial; for example, kinds of freedom that require reasons responsiveness (see, e.g., Fischer and Ravizza 1998 on this topic) might be somewhat problematic because one might think that the literature on things like situationism and confabulation suggests that we're less reasons responsive than we might have thought; but even if this is true, it seems pretty obvious that lots of our decisions are at least significantly influenced by our reasons. In any event, it seems to me that the *main* controversial subquestion of the which-kinds-of-freedom-do-we-have question is the libertarian question, that is, the question of whether we're L-free. If this is right, then it's a consequence of this paper that the metaphysical component of the problem of free will boils down largely, though perhaps not entirely, to the libertarian question. (Of course, one might think we've already got good reasons to doubt that we're L-free, but I think it can be argued that this is false.⁹)

Another consequence of this paper is that the answers to the what-is-free-will question and the which-kinds-of-freedom-do-we-have question might be very different. For instance, it may be that (a) human beings do possess L-freedom so that metaphysical libertarianism is true, but that (b) compatibilism is also true because the ordinary term 'free will' denotes some compatibilist kind of freedom and not L-freedom. There is no incompatibility between these two theses—indeed, there's not even any tension between them—and so there is no good reason for thinking that one shouldn't endorse them together.

Finally, it's worth noting that the point I have argued in this paper—that the compatibilism question and the what-is-free-will question are irrelevant to the do-we-have-free-will question—is a special case of a more general point, namely, that conceptual analysis is essentially irrelevant to metaphysics. I cannot argue this here, but the main idea should be clear: we cannot make any nontrivial progress toward discovering the nature of the world (or at any rate, the nonsemantic part of the world) by analyzing a concept, that is, by figuring out what some word means or should mean. Now, of course, if we already knew what some part of the world was like but didn't know what some folk term meant, then a conceptual analysis could tell us that certain about-the-world sentences were true; but in this case, we wouldn't be learning anything nontrivial about the nature of the (nonsemantic part of the) world.

Notes

I would like to thank Robert Kane, Timothy O'Connor, Carl Ginet, Kadri Vihvelin, John Martin Fischer, David Widerker, Robert Jones, Manyul Im, and David Pitt for commenting on earlier versions of this paper. Also, versions of this paper were read at the University of Oakland and the University of California, Riverside, and I would like to thank the members of both of those audiences for some very helpful feedback. Finally, the writing of this paper was partially supported by a fellowship from the National Endowment for the Humanities and also by a research grant from California State University, Los Angeles; I would like to express my gratitude to both of these institutions.

¹ One might wonder why philosophers are still worried about determinism given that the emergence of quantum mechanics has undermined any reasons we might have had for believing that determinism is true. There are two answers to this question. First, while it's true that we don't have any good reason to believe that determinism is true, we also don't have any good reason to believe that it's false, and so, for all we know, it might still be true. Second, even if determinism is false, there's a related thesis that one might still reasonably believe and that seems to generate a very similar problem for free will. The thesis in question—which might be called *FE-determinism*—is that there are no freedom-enhancing indeterminacies in human decision-making processes. Given that this might be true, one might wonder whether it's compatible with free will. But it turns out that free will is compatible with FE-determinism iff it's compatible with determinism. Thus, even if determinism is false, the compatibilism question might still be philosophically important.

² We can say that a person is libertarian free, or *L-free*, iff she makes at least some decisions that are such that (a) they are both undetermined (i.e., not causally determined by prior events) and appropriately nonrandom, and (b) the indeterminacy here is relevant to the appropriate nonrandomness in the sense that it generates it, or procures it, or enhances it, or increases it, or some such thing. Now, much needs to be said about what *appropriate nonrandomness* consists in, but we needn't bother with this here. Suffice it to say that the main

requirement is a kind of agent-involvedness that consists (most importantly, but not entirely) in the given agent having *authorship* and *control* over which option is chosen. See my 2004 and my forthcoming for more on this.

³ See Hume (1748) and Frankfurt (1971) for their notions of free will.

⁴ Parfit (1995) endorses a view similar to this in connection with the issue of personal identity.

⁵ Actually, let me say just a few words about the view that 'free will' is a kind term that picks out the sort of freedom that's inherent in normal human choices, whatever that turns out to be. To see how implausible this view is, notice how different 'free will' is from kind terms like 'water'. 'Water' denotes the stuff in our lakes and pipes and so on, whatever that stuff turns out to be. So if we discovered (in this world) that that stuff isn't H₂O, that it's really XYZ, then it would follow that water is XYZ; it certainly wouldn't follow that there is no such thing as water. But 'free will' doesn't work like this. Suppose we discovered that all our choices were controlled by Martians via remote control. If 'free will' were a kind term, it would follow that free will consists in being controlled by Martians. But, of course, that's wrong. If we discovered that Martians were controlling all our choices, that would be a discovery that we don't have free will.

⁶ This is perhaps not quite right. Consider, e.g., the view that free will is the ability to do otherwise; if this were right, then the question of whether free will is compatible with determinism would be controversial. But this is just because this definition is so thin and underdeveloped. In particular, we aren't told here what the ability to do otherwise is. Now, the question of what the ability to do otherwise is is obviously a controversial question. But, of course, it's not a compatibility question; it's a conceptual-analysis question and, indeed, on the view that free will is the ability to do otherwise, it's part of the what-is-free-will question. Thus, even if we have here a counterexample to the claim that there are no controversial subquestions of the which-kinds-of-freedom-are-compatible-with-determinism question, we don't have a good objection to my overall stance because we don't have a good objection to my claim that the compatibilism question reduces to the what-is-free-will question. For if we assume that free will is the ability to do otherwise, then the compatibilism question seems to turn on the question of what the ability to do otherwise is, and again, according to the view that free will is the ability to do otherwise, this question is a *part* of the what-is-free-will question.

⁷ For instance, Fischer (1994) thinks the argument shows that moral responsibility is compatible with determinism but not that free will is compatible with determinism.

⁸ All four of these strands are enormous, but let me give just a few references here. In connection with (i), aside from Hume 1748, see, e.g., Austin 1961, Chisholm 1964, Lehrer 1966, 1968, Aune 1967, and Berofsky 2002. In connection with (ii), see, e.g., Watson 1975, Dworkin 1988, Wolf 1990, Double 1991, Mele 1995, Bok 1998, and Haji 2002. In connection with (iii), see van Inwagen 1975, Wiggins 1973, Lamb 1977, Ginet 1980, and Kapitan 2002. In connection with (iv), see Frankfurt 1969 and Fischer 2002, and for some responses to the Frankfurt-case argument, see, e.g., Widerker 1995a, 1995b, Kane 1985, 1996, Ginet 1996, and Wyma 1997.

⁹ See, e.g., my 2004 and my forthcoming.

References

- Aune, B. 1967. Hypotheticals and 'can': Another look. *Analysis* 27:191–95.
- Austin, J. L. 1961. Ifs and cans. In *Philosophical papers*, ed. J. O. Urmson and G. Warnock, 153–80. Oxford: Clarendon Press.
- Balaguer, M. 2004. A coherent, naturalistic, and plausible formulation of libertarian free will. *Nous* 38:379–406.
- . Forthcoming. *Free will as an open scientific problem*. Cambridge, MA: MIT Press.
- Berofsky, B. 2002. Ifs, cans, and free will: The issues. In Kane 2002, 181–201.
- Bok, H. 1998. *Freedom and responsibility*. Princeton: Princeton University Press.
- Chisholm, R. 1964. J. L. Austin's philosophical papers. *Mind* 73:20–25.
- Double, R. 1991. *The non-reality of free will*. Oxford: Oxford University Press.
- Dworkin, G. 1988. *The theory and practice of autonomy*. Cambridge: Cambridge University Press.
- Fischer, J. M. 1994. *The metaphysics of free will: A study of control*, Oxford: Blackwell.
- . 2002. Frankfurt-type examples and semi-compatibilism. In Kane 2002, 281–308.
- Fischer, J. M., and M. Ravizza. 1998. *Responsibility and control: A theory of moral responsibility*. Cambridge: Cambridge University Press.
- Frankfurt, H. 1969. Alternate possibilities and moral responsibility. *Journal of Philosophy* 66:829–39.
- . 1971. Freedom of the will and the concept of a person. *Journal of Philosophy* 68:5–20.
- Ginet, C. 1980. The conditional analysis of freedom. In *Time and cause*, ed. P. van Inwagen, 171–86. Dordrecht: Reidel.
- . 1996. In defense of the principle of alternative possibilities: Why I don't find Frankfurt's argument convincing. *Philosophical Perspectives* 10:403–17.
- Haji, I. 2002. Compatibilist views of freedom and responsibility. In Kane 2002, 202–28.
- Hume, D. 1955 [1748]. *An inquiry concerning human understanding*. Indianapolis, IN: Bobbs-Merrill.
- James, W. 1956 [1884]. The dilemma of determinism. In *The will to believe; Human immortality*. New York: Dover Publications, 145–83.
- Kane, R. 1985. *Free will and values*. Albany: State University of New York Press.
- . 1996. *The significance of free will*. Oxford: Oxford University Press.
- . 2002. *The Oxford handbook of free will*. Oxford: Oxford University Press.
- Kant, I. 1927 [1788]. *The critique of practical reason*. Trans. T. K. Abbott. New York: Longman, Green, and Company.
- Kapitan, T. 2002. A master argument for incompatibilism? In Kane 2002, 127–57.
- Lamb, J. 1977. On a proof of incompatibilism. *Philosophical Review* 86:20–35.

- Lehrer, K. 1966. An empirical disproof of determinism. In *Freedom and determinism*, ed. K. Lehrer, 175–202. New York: Random House.
- . 1968. Cans without ifs. *Analysis* 29:29–32.
- Mele, A. 1995. *Autonomous agents: From self-control to autonomy*. New York: Oxford University Press.
- Nahmias, E., S. Morris, T. Nadelhoffer, and J. Turner. 2005. Surveying freedom: Folk intuitions about free will and moral responsibility. *Philosophical Psychology* 18:561–84.
- Nichols, S. 2006. Folk intuitions on free will. *Journal of Cognition and Culture* 6:57–86.
- Parfit, D. 1995. The unimportance of identity. In *Identity*, ed. H. Harris, 13–45. Oxford: Clarendon Press.
- van Inwagen, P. 1982 [1975]. The incompatibility of free will and determinism. In *Free will*, ed. G. Watson, 46–58. Oxford: Oxford University Press.
- Watson, G. 1975. Free agency. *Journal of Philosophy* 72:205–20.
- Widerker, D. 1995a. Libertarianism and Frankfurt's attack on the principle of alternative possibilities. *Philosophical Review* 104:247–61.
- . 1995b. Libertarian freedom and the avoidability of decisions. *Faith and Philosophy* 12:113–18.
- Wiggins, D. 1973. Towards a reasonable libertarianism. In *Essays on freedom of action*, ed. T. Honderich, 31–61. London: Routledge and Kegan Paul.
- Wolf, S. 1990. *Freedom within reason*. New York: Oxford University Press.
- Wyma, K. 1997. Moral responsibility and the leeway for action. *American Philosophical Quarterly* 34:57–70.