

Replies to McKenna, Pereboom, and Kane

Mark Balaguer

Published online: 25 March 2012
© Springer Science+Business Media B.V. 2012

The purpose of this essay is to respond to critiques of my recent book (*Free Will as an Open Scientific Problem*) put forward by Michael McKenna, Derk Pereboom, and Bob Kane in an Author-Meets-Critics session at the 2011 Pacific Division meeting of the APA. (Their critiques are also published in the current issue of the present journal, as McKenna (2012), Pereboom (2012), and Kane (2012).) Before I start, though, I would like to extend a heartfelt thank you to Michael, Derk, and Bob for the time and effort they put into the session and, of course, for their sage and useful comments. And I would also like to thank Joe Campbell for planning and organizing the session.

1 Replies to McKenna

1.1 McKenna's main objection

As I read McKenna, he raises five distinct objections, but he spends most of his time on one of them, so I will start there. To get at the real, underlying issue here, I need to provide a fair amount of background. To begin with, consider the following obviously central question:

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

If we ask what we really don't know here—that is, if we ask what we would need to know in order to fully answer the do-we-have-free-will question—it leads to a way of thinking of this question that's very useful. In particular, we can think of this question as reducing to (or collapsing into, or being subsumed by, or some such

M. Balaguer (✉)
Department of Philosophy, California State University, Los Angeles, CA, USA
e-mail: mbalagu@exchange.calstatela.edu

thing) two different questions, one semantic and one metaphysical (or one about conceptual analysis and one about the nature of human beings). The two questions are as follows:

The what-is-free-will question: What is free will? I.e., what is the correct analysis of the notion of free will?

and

The which-kinds-of-freedom-do-we-have question: Which kinds of freedom do humans have? I.e., do they have libertarian freedom?; and do they have Humean freedom?; and do they have Frankfurtian freedom?; and so on. (If we wanted to be more precise, we could formulate this question as asking which of the various kinds of “*freedom*” humans have—in scare quotes—because some or all of these things might fail to *be* free will, according to the correct answer to the what-is-free-will question. But I won’t always bother with being this precise here.)

I do not think McKenna has a problem with this way of thinking of the do-we-have-free-will question. His objections concern theses that I argue for later in chapter 2. The central thesis of chapter 2 is the following:

Central Thesis of Chapter 2: While the what-is-free-will question is obviously relevant to the do-we-have-free-will question in a certain way, it’s not relevant to that question in any metaphysically interesting way; in particular, it’s not relevant to the task of discovering the actual nature of human-decision-making processes, except in a trivial way. (And notice that if this is true, then it follows that the metaphysically interesting issue behind the do-we-have-free-will question is captured by the which-kinds-of-freedom-do-we-have question.)

My argument for this thesis is quite long—it takes up 15 pages of the book—and I will not try to summarize it here. But the central idea is as follows: If we already knew the answer to the which-kinds-of-freedom-do-we-have question—e.g., if we knew that humans had freedoms of kinds A, B, and C but did not have freedoms of kinds P, Q, and R—then we would not learn anything substantively new *about humans* by discovering that the term ‘free will’ is best thought of as referring to, say, C-freedom. This would, of course, tell us that human beings have free will, but given that we already knew which kinds of freedom humans have and do not have, this would not count as a genuinely new discovery about the actual nature of human decision-making processes, except in a trivial way.

It is important to note that my arguments here generalize. In other words, essentially the same arguments can be run in connection with other questions of semantics, or conceptual analysis—i.e., other questions that ask what the meaning of some expression is. Thus, if the arguments I give in chapter 2 are correct, then all questions of semantics, or conceptual analysis, are essentially irrelevant to questions about the actual nature of human beings.

This still is not the target of McKenna’s main objection, but he does voice a disagreement here. He says (2012):

I myself think that there is a more intimate connection between our concepts and the reality they are designed to articulate, and so we can learn about the way things actually are by attending to our concepts.

It is worth noting, however, that I would agree with McKenna that we can learn about the world by attending to concepts. For I acknowledge that the *articulation* of concepts—or of kinds of freedom—can lead to questions and discoveries about the nature of (the non-semantic part of) the world. What I claim is metaphysically irrelevant is the project of determining which (if any) of the various analyses of free will that we have articulated are *correct* analyses of the notion of free will.¹

In any event, after arguing that the what-is-free-will question is essentially irrelevant to metaphysical questions about the actual nature of human decision-making processes, I go on to argue that the compatibilism question is also essentially irrelevant to such questions. To see why this is the case, note first that just as we can think of the do-we-have-free-will question as reducing to the what-is-free-will question and the which-kinds-of-freedom-do-we-have question, so too we can think of the compatibilism question (i.e., the question of whether free will is compatible with determinism) as reducing to (or being subsumed by, or some such thing) the following two questions:

the what-is-free-will question; and
the which-kinds-of-freedom-are-compatible-with-determinism question.

Now, given this, it might seem obvious that the compatibilism question is relevant to questions about the nature of human decision-making processes. For the which-kinds-of-freedom-are-compatible-with-determinism question is obviously relevant to the which-kinds-of-freedom-do-we-have question, and the latter question is obviously about the decision-making processes of humans. In principle, this is right. So I would not claim that the compatibilism question is *in principle* irrelevant to questions about the nature of human beings. But I think it can be argued that the compatibilism question is, so to speak, for-all-practical-purposes irrelevant to such questions—at least given the current state of the debate. One way to argue for this conclusion—not yet the way that I argue for it in the book—would be to argue for the following claim:

(No-Controversial-Sub-Compatibility-Questions) There are no controversial or non-trivial subquestions of the which-kinds-of-freedom-are-compatible-with-determinism question. In other words, all of the various kinds of freedom (or “freedom”) that philosophers have articulated are either obviously

¹ In a footnote, McKenna says that because we are situated in the world, it is likely that our concepts bear certain relations to reality. This is undoubtedly true, but I do not think it is enough to undermine my position here. It is plausible to suppose that the fact that we actually possess certain kinds of free will is causally relevant to the fact that we have a concept of free will. But it is not very plausible to suppose that detailed, nuanced facts about the specific kinds of freedom we possess lead us to have a concept of free will that corresponds to these specific kinds of freedom. Thus, determining the exact nature of our concept of free will is not a reliable way of determining which kinds of freedom we have. For instance, and most notably in the present context, suppose we discovered that the ordinary concept of free will is (is not) a libertarian concept; this would not give us any good reason to think that we actually possess (do not possess) libertarian freedom. There is a lot more to say about this, but unfortunately, I cannot get into it here.

compatible with determinism or obviously incompatible with it. E.g., Humean freedom is obviously compatible with determinism; libertarian freedom is obviously incompatible with determinism; and so on.

If this thesis were true, then—given the current state of the debate and for all practical purposes—the compatibilism question would reduce to the what-is-free-will question. For in this case, the only controversial part of the compatibilism question would be the what-is-free-will question. And if this were true, then given my earlier conclusion (that the what-is-free-will question is essentially irrelevant to metaphysical questions about the actual nature of human decision-making processes), it would follow that the compatibilism question is also essentially irrelevant to such questions.

But, again, I do not argue the point in this way, because I do not think that (No-Controversial-Sub-Compatibility-Questions) is true. What I do think, though, is this:

While (No-Controversial-Sub-Compatibility-Questions) isn't strictly speaking true, the way in which it's false doesn't undermine my overall stance here; in particular, it remains true that (given the current state of the debate) *the only controversial parts of the compatibilism question are semantic questions.*

To bring out my argument for this, let's start with an example. Consider the analysis that takes free will to be the ability to do otherwise. This analysis generates the following compatibility question:

(Q) Is the ability to do otherwise compatible with determinism?

Since (Q) is controversial and non-obvious (and since it is a subquestion of the which-kinds-of-freedom-are-compatible-with-determinism question), it follows that (No-Controversial-Sub-Compatibility-Questions) is false. But it seems to me that the *reason* (Q) is controversial and non-obvious is that we still have some conceptual analysis to do. In particular, we have to figure out what the ability to do otherwise *is*. If we give this notion a libertarian analysis, it will be incompatible with determinism, whereas if we give it, say, a Humean analysis, it will not be. So the controversial question here is still a semantic question.

Now, of course, this is just an example, but it seems to me that the point here can be generalized. In particular, I endorse the following general claim:

Sweeping Claim: Consider all the different analyses of free will that have actually appeared in the mainstream literature. For any such analysis—call it F—if it is not obvious whether F-type freedom is compatible with determinism, then the reason it's not obvious is that there are non-obvious questions about how to understand one or more concepts at work in F. The reason is *not* that there is a non-obvious *logical* question (or a non-obvious pure compatibility question—more on what this means in a moment) about whether a sufficiently precisely understood kind of freedom—F-type freedom—is compatible with determinism.

This seems right to me, and if it is, then the compatibilism question reduces—for all practical purposes and relative to the current state of the debate—to semantic

questions, or questions of conceptual analysis. Thus, if I am right that semantic questions are essentially irrelevant to questions about the actual nature of human beings, then it follows that the compatibilism question is also essentially irrelevant to such questions—again, for all practical purposes and relative to the current state of the debate.

Before moving on to McKenna's main objection, it is worth pausing to note what the alternative to Sweeping Claim is. Now, one way to deny Sweeping Claim would be to maintain that the reason certain compatibility questions are non-obvious is that we do not have a sufficiently precise handle on what determinism amounts to; but I will ignore this option here, because (a) it is pretty implausible that this is really the source of the non-obviousness of any mainstream compatibility questions, and (b) even if it were, it would not matter in the present context, because this would favor my point of view, not McKenna's. Thus, I will assume that the main alternative to Sweeping Claim is the following:

Alternative to Sweeping Claim: There is at least one analysis of free will in the mainstream literature (or at least one kind of free will, call it *F-freedom*) such that (a) it's not obvious whether F-freedom is compatible with determinism; and (b) the source of the non-obviousness doesn't have anything to do with any questions of semantics or conceptual analysis—or to put the point a bit differently, we've got a sufficiently precise handle on what F-freedom and determinism amount to; and (c) there's a non-obvious *pure compatibility* question about whether F-freedom is compatible with determinism.

I am using 'pure compatibility question' here as a sort of catch-all phrase; it covers questions of a couple of different kinds that differ from one another in connection with the kinds of possibility that they ask about. The primary cases that I am thinking of are those where the relevant question is a *logical* question, in particular, a question about whether the conjunction of determinism and all the relevant definitions (i.e., the definitions of the terms at work in the formulations of F-freedom and determinism) logically entails that no one is F-free. But I want to allow that there could also be cases where the pure compatibility question was about *metaphysical* possibility, rather than logical or conceptual possibility.

In any event, we are now ready to discuss McKenna's main objection. He puts his worry here as follows (2012):

For most of the compatibilist proposals..., one can find a substantive point of contention as to whether the freedom...specified is compatible with determinism.

Before giving my main response to this objection, I would first like to voice a disagreement with the strength of McKenna's claim here. It seems to me that for just about all of the compatibilist proposals in the literature, *if we read them charitably*, then the given kinds of freedom are pretty obviously compatible with determinism. But this does not really matter in the present context, and it is not the main point I want to make in response to McKenna. The main point I want to make is the following: Given the above remarks, the important question here is not whether there are controversial compatibility questions of the kind that McKenna has in

mind; the important question is whether Sweeping Claim is true of these compatibility questions. And it seems to me that the cases McKenna discusses actually fit perfectly with Sweeping Claim. Consider, e.g., the issue regarding Frankfurt-freedom discussed by Velleman (1992). The issue here is how the notion of identification (and the notion of an agent *doing* something) are to be understood. This sounds to me like a semantic question, not a logical question or any other kind of pure compatibility question. In other words, what we do *not* have here are (a) a sufficiently precise definition of identification, or of agential action, and (b) a non-obvious pure compatibility question about whether the given (sufficiently precisely defined) notion is compatible with determinism. Notice, for instance, that if we understand identification in the way that Frankfurt did in his original 1971 paper, then it is pretty obviously compatible with determinism; and if we understand agential action in the way that Velleman does, then it is obviously compatible with determinism as well; and if we understand this sort of action as involving libertarian agent causation, then it is obviously *not* compatible with determinism; and so on.

(By the way, for whatever it's worth, Velleman seems to agree with me that what is at issue here is a question of conceptual analysis. He says (p. 466) that his aim is "to explain what we ordinarily mean when we call something an action".)

Similar remarks can be made about McKenna's other examples. Consider, e.g., the worry that McKenna himself raises for Fischer-Ravizza-freedom in his (2001). The issue here is how the term 'same mechanism' is to be understood. Once again, this seems to be a semantic question. And likewise for Vihvelin's (2004) view. I agree with McKenna (and I suspect that Vihvelin would *disagree* with this) that it is not obvious whether the dispositions that Vihvelin has in mind are compatible with determinism; but it seems to me that the *reason* for this is that it is not obvious how these dispositions are to be understood. What we do not have here are (a) a perfectly precise statement of what the relevant dispositions consist in and (b) a non-obvious logical question, or pure compatibility question, about whether such dispositions are compatible with determinism.

Finally, it is worth remembering that I do not really *need* Sweeping Claim. For, again, when I say that Sweeping Claim is true, I do not mean to be making an in-principle claim; I am just making a claim about the analyses of free will that are currently on offer in the mainstream literature. If tomorrow someone came up with an analysis of free will that gave rise to a non-obvious compatibility question, and if the ultimate source of the non-obviousness was a non-obvious *logical* question, as opposed to non-obvious semantic questions, that would not bother me. I would simply admit that the given compatibility question was relevant to questions about the metaphysical nature of human beings. And, likewise, if McKenna convinced me that Sweeping Claim is false—i.e., that there is a non-obvious pure compatibility question associated with one of the kinds of freedom currently on offer in the literature—I would react in the same way. None of this would undermine my overall view because the *reason* these compatibility questions would be relevant to questions about the actual nature of human beings (e.g., to the do-we-have-free-will question) is that they would be relevant to the which-kinds-of-freedom-do-we-have question. So this would not undermine my claim that the which-kinds-of-freedom-do-we-have question is the metaphysically interesting question behind the

do-we-have-free-will question, and it would not undermine my claim that semantic questions are essentially irrelevant to metaphysical questions about the actual nature of human decision-making processes.²

1.2 McKenna's other objections

Let me move on now to four other objections that McKenna raises...

1. Consider the following meta-semantic question:

Meta-semantic question: What sorts of facts determine what free will is, or what 'free will' means, or which analyses of free will are correct?

McKenna says that I think that the relevant facts here are at least partially facts about the usage and intentions of ordinary folk, and he then argues against this view. But I actually never committed to the view that McKenna ascribes to me here; indeed, in the book, I did not commit to *any* answer to the meta-semantic question. Moreover, if I were to take a stand here, I would want to say (with a few caveats or provisos about what exactly I meant to be saying) that there is no determinate fact of the matter about the answer to the meta-semantic question. Different philosophers seem to want to answer it in different ways, and I think this leads to people talking past each other sometimes. But I do not think we need to *figure out* the answer to the meta-semantic question. I think we just need to be clear about what we are doing. If you want to try to figure out what ordinary folk mean by 'free will', then be clear that that is what you are doing; and if you want to try to figure out what mainstream philosophers mean, then be clear about that; and if you want to try to figure out what kinds of free will are required for moral responsibility, then be clear about that; and so on.

It is also worth noting here that there is a clear sense in which my book actually pushes the idea that we should not be very concerned with ordinary language. For one of my central contentions is that we can clarify things by moving away from folk expressions like 'free will' and replacing them with terms of art with stipulated definitions—e.g., 'libertarian freedom', 'Humean freedom', 'reasons-responsive freedom', and so on. In particular, I think we can clarify all the different issues by lining up all the different kinds of freedom that we might care about and asking various questions about them—e.g., 'Which of these kinds of freedom do we have?'; 'Which of them are compatible with determinism?'; 'Which of them are required for moral responsibility?'; 'Which of them are worth wanting?'; and so on.

2. McKenna accuses me of thinking that there's a single, unique problem of free will. But I actually do not think that. I guess I sometimes use the phrase '*the* problem of free will' to denote a certain problem, but as the above remarks make clear, I think it is worth asking various different questions about all the different kinds of freedom (or "freedom") that people might care about.

² Consider the view that 'free will' denotes the kind of freedom that human beings actually have, i.e., the kind of freedom that's actually inherent in choices of the kind that we ordinarily call free, whatever this kind of freedom turns out to be. You might think that this analysis gives us a counterexample to Sweeping Claim. I do not have the space to discuss this worry here, but see pp. 48–50 of my book.

3. At one point in the book, in arguing that the jillions of little decisions we make, taken together, are just as important (if not more important) than the handful of big, life-changing decisions we make, I say this (p. 112):

If I weren't in control of whether I had eggs or cereal for breakfast, or whether I exercised before going to work, or whether I went into the office or worked from home..., then I wouldn't have control over my *life*.

McKenna responds by pointing out that even if we did not have libertarian-style control in these cases, we would still have various compatibilist kinds of control. He is right about that, and in fact, I did not mean to deny this. The above remark is perhaps not as clear as it could have been. Here's what I meant: If we do not have control of a given kind *K* over any of the little decisions that we make, then we do not have kind-*K* control over our lives. I was not arguing in this passage that libertarian control is more important than other kinds of control. Rather, I was arguing that having control (of whatever kind you are interested in) over the jillions of little decisions you make is just as important as having control (of the relevant kind) over the handful of big decisions you make.

(Earlier in chapter 3 (see especially pp. 96–106), I also argued that if we do not have libertarian control (of the sort I discuss in the book) over decisions of a certain kind (I will call them *torn* decisions, and I will define them below), then we have less overall control over those decisions; but I was talking there about torn decisions only, not all decisions. And, again, in the passage quoted above, I did not mean to be saying anything about the relative importance of libertarian control and other kinds of control; I was talking only about the relative importance of little decisions and big decisions.)

4. I argue in the book that what might be called the *which-kinds-of-freedom-are-required-for-moral-responsibility question* is essentially irrelevant to questions about the actual nature of human-decision-making processes (in a nutshell, this is because the question here about moral responsibility is essentially a *semantic* question). McKenna responds as follows:

...in learning what sort of freedom is required for moral responsibility, and in establishing that at least some persons are morally responsible for what they do, we would learn something substantive about normally functioning human deliberative processes.

I actually responded to a worry just like this in my book (mainly on pp. 42–43, but see also 39–41 and 137–141). Unfortunately, I do not have the space to rehearse this response here.

2 Replies to Pereboom

Pereboom's central objection to my view—indeed, to any event-causal libertarian view—is what he calls the disappearing-agent objection. To get at this objection, we need a bit of background. To begin with, the central question that is at issue here is

the question of whether human beings have libertarian free will, or *L-freedom*. We can define *L-freedom* as follows:

A person is *L-free* iff she makes at least some decisions that are such that (a) they're both undetermined and appropriately non-random, and (b) the indeterminacy is relevant to the appropriate non-randomness in the sense that it *generates* the non-randomness, or *procures* it, or *enhances* it, or *increases* it, or something along these lines. (A lot needs to be said about what appropriate non-randomness consists in, but for our purposes here, all that matters is that it's a kind of agent-involvedness and that, at the very least, it involves authorship and control—i.e., it requires it to be the case that the agent authors and controls the given decision.)

There is a very old argument for thinking that human beings do *not* have *L-freedom* that can be put like this:

Traditional anti-libertarian argument: If an event is not causally determined, then it's random in a certain sense. Thus, the insertion of an undetermined event into a decision-making process could never increase the non-randomness of that process, and so human beings could not be *L-free*.

One of the central arguments of my book turns this traditional argument on its head. I argue that there is a certain category of our decisions (I call them *torn decisions*, and I will characterize them in a moment) for which the following is true: If they are undetermined in the right way, then they are also appropriately non-random and *L-free*. Notice that I am not just saying that our torn decisions *could be* both undetermined and non-random; I am saying that if they are undetermined in the right way, then they *are* appropriately non-random and *L-free*, so that the question of whether our torn decisions are *L-free* (and also, I think, whether *we* are *L-free*) just reduces to the question of whether these decisions are undetermined in the right way.

To make this more precise, I need to say what a torn decision is, and I need to specify what the relevant sort of indeterminism is. So, first, let me say this:

A *torn decision* is a decision in which the person in question has reasons for multiple options, feels torn as to which option is best, and decides without resolving the conflict, i.e., decides while feeling torn.

I think we make decisions like this several times a day about things like whether to have eggs or cereal for breakfast, or whether to work out before leaving for work, or whatever. But we can also make torn decisions in connection with big life-changing decisions; e.g., you might have a good job offer in a bad city, and you might have a deadline that forces you to decide while feeling utterly torn. (Torn decisions are obviously a lot like Kane's self-forming actions, or SFAs. But there are a few differences. Note, in particular, that unlike SFAs, torn decisions are not defined as being undetermined. They are defined in terms of their phenomenology. Thus, we know from experience that we do make torn decisions, and it is an empirical question whether any of these decisions are undetermined in the right way. For more on the differences between torn decisions and SFAs, see pp. 73–75 of my book.)

Next, let me define the relevant sort of indeterminism, i.e., the sort that is needed for an ordinary torn decision to be fully L-free. We can do this as follows:

A torn decision is *wholly undetermined* at the moment of choice—or, as I'll also say, *TDW-undetermined*—iff the actual, objective moment-of-choice probabilities of the various reasons-based tied-for-best options being chosen match the reasons-based probabilities (or the phenomenological probabilities), so that these moment-of-choice probabilities are all roughly even, given the complete state of the world and all the laws of nature, and the choice occurs without any further causal input, i.e., without anything else being significantly causally relevant to which option is chosen.

Given these definitions, the thesis that I argue for in my book can be put very succinctly:

Central Thesis of Chapter 3: If our torn decisions are undetermined in the right way—i.e., if they're wholly undetermined, or TDW-undetermined—then we author and control them, and they're appropriately non-random and L-free.

My argument for this thesis takes up a large chunk of the book, so I obviously cannot rehearse the whole argument here. Nonetheless, I need to say a few words about it. To bring out the central idea behind the argument, let me assume that a weak, token-token mind-brain identity theory is true, so that ordinary human decisions are neural events. I do not actually need this assumption, but it makes the argument run more smoothly. In any event, given this background assumption, let's look at an example of a torn decision. Suppose that Ralph has to choose between two options, O and P, and suppose that he makes a torn decision to go with O rather than P. The details do not matter; option O could be something important like a new job, or it could be a bowl of lobster bisque. All that matters is that Ralph makes a conscious torn decision to go with option O. Given this, if we assume that Ralph's decision was TDW-undetermined, we get the following results:

- (A) Ralph's choice was conscious, intentional, and purposeful, with an actish phenomenology—in short, it *was* a Ralph-consciously-choosing event, or a Ralph-consciously-doing event (we actually know all of this independently of whether the choice was TDW-undetermined); and
- (B) the choice flowed out of Ralph's conscious reasons and thought in a nondeterministically event-causal way; and
- (C) nothing external to Ralph's conscious reasons and thought had any significant causal influence (after he moved into a torn state and was about to make his decision) over how he chose, so that the conscious choice itself *was* the event that settled which option was chosen. (If you like, we can put it this way: The conscious choice itself *was* the undetermined physical event that settled which option was chosen.)

My claim is that given (A)–(C), it makes good sense to say that Ralph authored and controlled the decision. It seems to make sense in this scenario to say that *(i) Ralph*

did it, and (ii) nothing made him do it. And, intuitively, this seems to be sufficient for authorship and control.

Pereboom disagrees. He thinks that in this case, we do *not* get the result that Ralph “did it”, or that Ralph authored and controlled the decision. We can put his argument here in the following way:

Pereboom’s disappearing-agent objection: Full-blown authorship and control require the *agent* to settle which option is chosen. But if a torn decision is TDW-undetermined, then the agent *doesn’t* settle it. Indeed, *nothing* settles it. Pereboom doesn’t put it quite this way, but his idea seems to be that if a torn decision is TDW-undetermined, then nothing *causes* the given option to be chosen—or nothing *causally settles* which option is chosen—and so the choice of the given option *just happens*, and so the person in question doesn’t author or control the decision.

If we apply this objection to Ralph’s decision, the central claim seems to be that in order for it to be the case that Ralph authored and controlled his decision, it needs to be the case that he caused option O to be chosen. Or at any rate, Pereboom thinks this is so if ‘cause’ is understood as expressing a make-happen relation, which I am willing to grant for the sake of argument. In short, then, Pereboom’s objection is based on the following claim:

(AC) Authorship and control require agent causation (where agent causation involves the agent *causally settling* which option is chosen).

(Actually, Pereboom seems to endorse something even stronger than (AC). He seems to think that if Ralph did not cause option O to be chosen, then it could not be the case that he was the subject of a conscious, actish-feeling, purposeful, intentional decision. This is hard for me to understand. Presumably, he would not deny that the event was conscious or actish-feeling, so he must be saying that it could not have been an intentional, purposeful decision. This strikes me as a really strong claim; it seems to me that we know via first-hand experience that we *do* make intentional, purposeful decisions, regardless of how they are caused. But I do not want to pursue this here because Pereboom does not actually need this ultra-strong claim; all he needs is (AC), for this already gives him the result that if Ralph did not cause option O to be chosen, then he did not author and control the decision, and the decision was not L-free. So let me focus on (AC).)

How can I respond to Pereboom’s objection? Well, the most obvious response would be to argue that Pereboom is simply mistaken about (AC), i.e., that agent causation is in fact *not* required for authorship and control. I will say a few words to motivate this stance below, but it is not the main response I want to make to Pereboom’s objection. The main response I want to make is that I do not need to claim that Pereboom is mistaken about (AC). To bring this point out, let me start by distinguishing two different kinds of control—*DP-control* and *MB-control* (the names, if you’re wondering, come from our initials)—where the former requires agent causation (or something like it) and the latter does not. More precisely, MB-control is a kind of control that applies to ordinary torn decisions if they are TDW-undetermined; i.e., it applies to torn decisions like Ralph’s, where the agent makes a

conscious decision with an actish phenomenology and which option is chosen is not significantly causally influenced (at the moment of choice) by anything external to the agent's conscious reasons and thought, so that the conscious choice itself *is* the event that settles which option is chosen. Of course, this still does not give us a precise definition of MB-control, but I do not want to try to do any better than this. In the book, I purposely avoided giving a necessary-and-sufficient-condition-style conceptual analysis of control, and for a variety of reasons, I still want to avoid doing this. But for the sake of argument, let's pretend that we've got hold of the kind of control I am waving my hands at here, and let's call it "MB-control". And, likewise, let's pretend that we've got hold of a unique kind of control that requires agent causation, and let's call it "DP-control". (We can also pretend that we've got hold of two different kinds of authorship—*MB-authorship* and *DP-authorship*—but in what follows, I will be focusing more on control than on authorship.)

Given that we've got these two analyses of control, or these kinds of control, one question we might ask is the following:

The what-is-control question: What is control? I.e., which of the various kinds of control that we find in the literature are *genuine* kinds of control? Or if you'd rather, which of these kinds of control provide correct analyses of the concept of control? E.g., does DP-control provide a correct analysis of control?; does MB-control?; does neither?; do both?

This, however, is a *semantic* question. Thus, if the arguments in chapter 2 of my book are correct, then the what-is-control question is essentially irrelevant to questions about the metaphysical nature of human decision-making processes. If we assume that Ralph's decision was TDW-undetermined, then it follows that Ralph MB-controlled his decision but did not DP-control it; these are the metaphysical facts *about Ralph*; answering the what-is-control question would not give us any genuinely new information about Ralph; it might tell us the truth value of the English sentence 'Ralph controlled his decision', but since we already know that he MB-controlled the decision but did not DP-control it, this would not count as a genuinely new discovery *about Ralph*, except in a trivial way.

But even if this is right, you might think I need to motivate the conceptual claim that MB-control is real control, or that MB-control provides a correct analysis of the concept of control. You might think I need to do this in order to obtain the result that if our torn decisions are TDW-undetermined, then we *control* them. But I actually do not need this result. I do not need to claim that if our torn decisions are TDW-undetermined, then they are authored and controlled by us and L-free in the only senses of these terms that anyone might care about, or in the senses of these terms that philosophers have traditionally cared about. All I need is this:

(*) If our torn decisions are TDW-undetermined, then they're authored and controlled by us and appropriately non-random and L-free in interesting and important ways that are worth wanting and worth arguing for and that libertarians can hold up and say, "This gives us a noteworthy kind of libertarian free will." (Actually, what I really need here is a bit weaker than this; all I really need is that if our torn decisions are TDW-undetermined, then

we “author” and “control” them and so on in ways that are interesting, important, and so on; for on the view I’m suggesting, it doesn’t matter whether MB-authorship and MB-control provide correct analyses of the concepts of authorship and control, and so it wouldn’t matter if they failed to be genuine referents of the English terms ‘authorship’ and ‘control’. But I won’t bother with this complication here.)

Now, don’t take me to be saying more than I am here. I am not saying that libertarians can define authorship and control and L-freedom *however they want to*; they cannot just define these terms in ridiculously weak ways and then claim victory. I do not need to argue that the kind of L-freedom I have articulated—the kind we get if our torn decisions are TDW-undetermined (i.e., the kind that involves MB-authorship, MB-control, and so on)—is the one and only kind of L-freedom that anyone might care about. But I do need it to be the case that this kind of L-freedom is interesting, worth wanting, worth arguing for, and so on. In other words, I need (*).

But I think the arguments in my book do motivate (*). Let’s return to Ralph’s decision. If it is TDW-undetermined, then, again, we have the following results:

(A) Ralph’s choice was conscious, intentional, and purposeful, with an actish phenomenology—in short, it *was* a Ralph-consciously-choosing event, or a Ralph-consciously-doing event; and

(B) the choice flowed out of Ralph’s conscious reasons and thought in a nondeterministically event-causal way; and

(C) nothing external to Ralph’s conscious reasons and thought had any significant causal influence (after he moved into a torn state and was about to make his decision) over how he chose, so that the conscious choice itself *was* the event that settled which option was chosen.

It seems to me that even if (A)–(C) are not enough to give us every kind of L-freedom you might have wanted, they are clearly enough to give us *one important kind* of L-freedom—a kind that libertarians can hang their hats on and that is worth wanting and arguing for and so on. After all, in this scenario, the event that settles which option is chosen *is* the conscious decision—i.e., it is the event with a me-consciously-choosing-now phenomenology. One of the central claims in Pereboom’s disappearing-agent objection is that authorship and control require the *agent* to settle the matter. I am OK with that way of putting things. But it seems to me that if the event that settles the matter is the *agent’s conscious decision*, then, at the very least, there is a *sense* in which the agent does settle it. There might be other senses—most notably, agent-causal senses—in which the agent does *not* settle it; but, again, all I need is that there is one interesting, important sense in which the agent does settle it. And it seems to me that if the event that settles which option is chosen is the agent’s conscious choice, then that clearly gives us a sense in which the agent settles it.

In the book, I say a lot more about this. In particular, I say a lot more to motivate the idea that the sort of L-freedom that I am describing here is worth wanting. There are a few worries you might have about this. For instance, you might think that torn

decisions are trivial and that there is no good reason to care very much about them. Or you might worry that the kind of indeterminacy that I describe is not worth wanting because it is just a mechanism that helps us get unstuck when we are torn (e.g., you might think that our torn decisions might as well be decided by coin tosses). Or you might worry that the sort of L-freedom I am describing is not needed for moral responsibility, or that it is not sufficient for moral responsibility. I respond to these worries in the book, as well as a few others, and in the process, I argue that the sort of L-freedom I am describing here is worth wanting. Unfortunately, though, I do not have the space to rehearse these arguments here.

So that's my main response to Pereboom: I do not need to argue that MB-control gives us a correct analysis of the concept of control, or that the kind of L-freedom I describe in my book is the only kind of L-freedom that anyone might care about; all I need to argue is that if our torn decisions are TDW-undetermined, then they are L-free and appropriately non-random and authored and controlled by us (or "authored" and "controlled" by us, as the case may be) in ways that are interesting and important and worth wanting and that libertarians can hold up and say, "This gives us a noteworthy kind of libertarian free will." That's all I need, and I think I have delivered it.

But there is also a second point I would like to make in response to Pereboom's objection: Even though I do not need this result here, I think there is something to be said for the claim that MB-control gives us a better analysis of the concept of control than DP-control does. In other words, I think there is something to be said for the claim that agent causation (of the kind that involves causal settling) is in fact *not* needed for control—or, for that matter, for authorship. This is tricky business because the relevant notions of authorship and control are not really folk concepts, and given this, I think it is pretty hard to motivate the idea that there is an objective fact of the matter about how they should be analyzed. One thing we could do here is poll the intuitions of philosophers, but that's pretty obviously going to give us a split decision. I have a clear intuition that agent causation is not needed for authorship or control, but others (e.g., Pereboom and Tim O'Connor) have the opposite intuition. I will not attempt to settle this question here, but I would just like to offer one quick argument against the idea that authorship and control should be thought of as requiring agent causation. Let me put the argument in the form of a challenge to advocates of agent-causal analyses of authorship and control. The challenge is to say what exactly is to be *gained* by requiring agent causation. On the view I have in mind, we say that Ralph authored and controlled his decision (if it was TDW-undetermined) because (roughly) the event that settled which option was chosen was the conscious decision itself. If you demand that Ralph caused option O to be chosen (or that he causally settled which option was chosen), then it seems to me that you have gained nothing; you have simply moved everything back a step. For now there is a second event, on top of the conscious decision—namely, the event of Ralph agent-causing the decision—and we can ask the very same question about this event that Pereboom wants to ask about the conscious decision; that is, we can ask what caused the agent-causal event to occur. And, of course, the agent-causal response is going to be that *nothing* caused it to occur. But then how has anything

been gained? And if nothing's been gained, then why should we think of authorship and control as requiring agent causation?

Finally, let me say a few words about the event-causal-vs-non-causal-vs-hybrid issue that Pereboom brings up. If you like, you can think of the libertarian view I describe as a hybrid view, for while it involves some core event-causal theses (e.g., that our torn decisions are probabilistically caused by agent-involving events), it also has some non-causal threads in it. But as I point out in the book, you might think there are non-causal threads in all sensible event-causal views. For you might think that probabilistic causation is just deterministic causation of probabilistic states that then issue by chance (i.e., in an uncaused way) into one possible outcome or another (see, e.g., Hausman 1998). If we take this line, then the natural event-causal view of, e.g., Ralph's decision will be that (a) Ralph's reasons deterministically caused it to be the case that the probabilities of his two tied-for-best options being chosen were 0.5 each and (b) beyond this, which option was chosen was uncaused.

3 Replies to Kane

3.1 Preliminaries

Let me start by saying that I have no problem with Kane's claim that plural voluntary control (i.e., PVC) is needed for full-blown L-freedom. I argue in the book that if our ordinary torn decisions are TDW-undetermined, then they're also appropriately non-random and L-free, and I think this will involve all of the components of PVC. He lists these components in the following passage (2012):

Agents have plural voluntary control over a set of options (e.g., choices) when they are able to bring about whichever of the options or choices they will to bring about, when they will to do so, for the reasons they will to do so, on purpose, rather than accidentally or inadvertently or by mistake, without being coerced or compelled in doing so or willing to do so or otherwise controlled in doing so or willing to do so by other agents or mechanisms.

Again, I have no problem with this. What Kane and I disagree about is what is needed in order to *procure* PVC and appropriate non-randomness and L-freedom. That is, we disagree about what our torn decisions need to be like in order to be L-free. I think that if they are TDW-undetermined, then they are also fully L-free. In other words, and somewhat roughly, I think TDW-indeterminism is a sufficient condition for L-freedom. (The reason this is "somewhat rough" is that I do not commit to the claim that *any* torn decision that was TDW-undetermined would be L-free, even if, say, aliens had planted reasons in the agent's head just prior to choice; all I commit to is the weaker claim that if *our actual, ordinary* torn decisions are TDW-undetermined, then they are also L-free; but I will ignore this complication here and write as if I endorse the claim that TDW-indeterminism is sufficient for L-freedom.) In any event, Kane disagrees with me here; he thinks that

if a torn decision is to be fully L-free, then at the very least, the following conditions need to be satisfied:

- (i) The decision needs to be the end product of an indeterministic goal-directed cognitive process (or volitional stream, or effort of will), where the indeterminism involved in this process is a hindrance or an obstacle to be overcome; and (ii) there need to be at least two different volitional streams of this kind occurring in the brain during the decision-making process—streams that are aiming toward different tied-for-best options (or at least different live options)—and the agent needs to have a kind of one-way (i.e., non-plural) guidance control over each of the relevant volitional streams.

Let me refer to these conditions as the *Kane-conditions*. Again, Kane thinks these conditions are necessary for L-freedom, and so he rejects my claim that TDW-indeterminism is sufficient for L-freedom (in ordinary torn decisions).

Some of Kane's remarks seem to suggest that he thinks the Kane-conditions are strictly stronger than the condition of TDW-indeterminism, so that to satisfy the Kane-conditions, a torn decision needs to be TDW-undetermined and also satisfy some other conditions. But I do not think Kane really believes this. I think that (a) the Kane-conditions require some things that TDW-indeterminism does not require, and (b) TDW-indeterminism requires some things that the Kane-conditions do not require, so that neither TDW-indeterminism nor the Kane-conditions are strictly stronger. And I think Kane would agree with me on this point.

Nevertheless, even if TDW-indeterminism is not strictly weaker than the Kane-conditions, it seems to me—and I think Kane would agree with me about this as well—that, all things considered, TDW-indeterminism is an easier condition to satisfy than the Kane-conditions. If this is right, and if I am also right that TDW-indeterminism is sufficient for L-freedom, then I think it is a point in favor of my view. For I think it should be an aim of libertarians to locate the most minimal, or the easiest-to-satisfy, sufficient condition for L-freedom.

In any event, let's get a bit clearer on the differences between TDW-indeterminism and the Kane-conditions. And let me start by discussing the ways in which the Kane-conditions go beyond TDW-indeterminism. There are two main points that need to be discussed here. First, in order for a torn decision to satisfy the Kane-conditions, it needs to be the output of a rather substantive cognitive process. Now, of course, if a torn decision is TDW-undetermined, it will also be the output of a cognitive process. In particular, the agent has to have reasons for both options, she has to feel torn, and she has to decide while felling torn; moreover, the moment-of-choice probabilities have to match the phenomenological probabilities, and which option is chosen has to be free of any significant moment-of-choice causal influences from external factors; finally, all of these requirements more or less guarantee that the choice will flow out of the agent's reasons in a non-deterministic event-causal way. So, again, in order for a torn decision to be TDW-undetermined, it needs to be the output of a certain kind of cognitive process. But Kane requires that the process be *thicker*, or *meatier*, in a certain way. Now, it is not entirely clear how exactly to articulate the required "meatiness" here, but it seems to involve at least a certain kind of *trying*, or *effort*. (Kane also says in this connection that the

agent needs to be involved in a certain kind of *activity*—in particular, a goal-directed, guidance-controlled, cognitive activity—but I think that TDW-indeterminism delivers this result as well, so I will ignore this requirement here.) In any event, Kane also requires that there be *multiple* “meaty” processes of the relevant kind, so that in order for a torn decision to satisfy the Kane-conditions, the agent needs to have been involved in two different tryings, or efforts; she needs to have been trying to achieve two different (presumably incompatible) goals.

The second way in which the Kane-conditions seem to go beyond TDW-indeterminism is as follows: The indeterminism in the decision-making process has to be an *obstacle* to be overcome. This strikes me as an odd way of talking. It is suggestive of the idea that the indeterminism is a positive thing or event. On my way of thinking, the relevant indeterminism consists in a certain kind of *lack*; to say that a torn decision was TDW-undetermined is to say that it lacked certain kinds of causal influences. It would make more sense to me if Kane said that the obstacle to be overcome was the *opposing volitional stream* (together, perhaps, with the fact that the opposing stream could have won out because it was not causally determined in advance which stream would win out). I will assume that this is in fact what Kane means when he speaks of indeterminacy being an obstacle, although I might be wrong about this.

If this is right, then the two main ways in which the Kane-conditions go beyond TDW-indeterminism can be combined into one, and we can capture them both by saying this: The Kane-conditions require it to be the case that prior to the moment of conscious choice, the agent was involved in multiple “meaty” processes (or volitional streams, or tryings) that were competing with one another and creating noise for one another.

Now that we know how the Kane-conditions go beyond TDW-indeterminism, let’s ask the opposite question: How does TDW-indeterminism go beyond the Kane-conditions? In a nutshell, the answer is that in order for a torn decision to be TDW-undetermined, it needs to be the case that at the moment of choice (i.e., after the person is in a torn state and about to make a torn decision), nothing external to the agent’s conscious reasons and thought comes in and significantly causally influences which option is chosen. This does not seem to be required by the Kane-conditions.

In what follows, I will argue that (a) this no-external-causal-influence condition is in fact *needed* for full-blown L-freedom, and (b) the sorts of “meaty” cognitive processes (or volitional streams, or tryings) that Kane requires are *not* needed for full-blown L-freedom. Finally, at the end, I will respond to Kane’s luck-based argument for the claim that TDW-indeterminism is not sufficient for L-freedom.

3.2 Moment-of-choice external causal influences are freedom damaging

I want to argue in this section that in order for a torn decision to be full-blown L-free, it needs to be the case that at the moment of choice—i.e., after the person is in a torn state and about to make a torn decision—nothing external to the agent’s conscious reasons and thought comes in and significantly causally influences which option is chosen. There are two quick ways to appreciate this point. First, the no-significant-external-causal-influence condition is precisely what is needed to

give us the result that the conscious choice itself is the event that settles which option is chosen. Suppose that just prior to choice (after the person is in a torn state and so on), a nonconscious brain event occurs that significantly causally influences which option is chosen. Then we would *not* get the result that the conscious choice *is* the event that settles the matter. For the prior-to-choice nonconscious brain event would at least partially settle the matter. And given this, I think we would not get the result that the decision is fully L-free. And so it seems to me that we really do need the no-significant-external-causal-influence condition.

Second, suppose I make a torn decision and that, at the moment of choice, some factor that is completely external to my conscious reasons and thought comes in and significantly causally influences which option is chosen. Then, simplifying a bit, the causal factor would either be (a) completely external to my psychology (i.e., to my total mental state, or some such thing) or (b) some sort of subconscious mental state or event or process. But in both of these scenarios, we get the result that the causal factor damages or diminishes the L-freedom inherent in the decision—i.e., it diminishes the degree of agential control, or appropriate non-randomness, or something along these lines. In connection with (a), this is more or less obvious, since the causal factor is completely external to my psychology. In connection with (b), it is perhaps a bit less obvious, but I argue in the book (pp. 100–104) that, at the very least, there is a decrease here in the amount of *active, conscious control* that the agent has. I cannot run through the argument for this here, but the result is hardly surprising; given that the causal factor is unconscious (or that it exerts its causal influence unconsciously), it seems pretty hard to deny that it decreases the amount of active, conscious control that the agent has.

(It is important to note here that on my view, there can be *degrees* of causal influence by external factors; thus, the L-freedom of our torn decisions can be undermined to a greater or lesser extent, and so there can be degrees of L-freedom. Thus, I'm not saying here that if which option is chosen is causally influenced at all by external factors then it is not L-free; all I am saying is that for a torn decision to be *fully* L-free, it needs to be TDW-undetermined and, hence, which option is chosen cannot be significantly causally influenced, at the moment of choice, by external factors.)

3.3 “Meaty” cognitive processes or “tryings” are not needed for full-blown L-freedom

In order for a torn decision to satisfy the Kane-conditions, it needs to be the case that prior to the moment of conscious choice, the agent was involved in multiple “meaty” cognitive processes (or volitional streams, or tryings) that were competing with one another and creating noise for one another. But it is hard to see why this is needed for L-freedom. Suppose that I am presented with a choice between A and B, think about it for a moment, feel torn, and with no further ado, choose A in a TDW-undetermined torn-decision sort of way. And now suppose that essentially the same thing happens to Kane except that he becomes more tortured by the need to decide between A and B and, thus, engages in more substantive cognitive processes and efforts of will than I do, and suppose that these processes end up competing with

one another and creating noise for one another in the way that Kane has in mind. I do not see how this is supposed to increase the L-freedom inherent in his choice. In order for a torn decision of mine to be L-free, various things are needed. E.g., it needs to be the case that *I* make the decision and that nothing external to me makes me choose in the way that I do. But it does not need to be the case that I *agonize* over the decision. I suppose Kane thinks that if I agonize over it, then that will add to it being the case that *I* make the decision. But I don't see why. I think it can be *me* who makes the decision even if I do *not* agonize over it. Or more to the point, I think the decision can be fully L-free even if I do not engage in "meaty" cognitive processes of the kind that Kane has in mind.

Indeed, we can think of the arguments in chapter 3 of my book as establishing this point. I argue there that if our ordinary torn decisions are TDW-undetermined, then they are appropriately non-random and L-free. Thus, since TDW-indeterminism does not require agents to engage in "meaty" cognitive processes of the kind that Kane has in mind, it follows that if my argument is cogent, then these kinds of "meaty" cognitive processes are not needed for L-freedom.

But, of course, Kane does not think that the argument in chapter 3 of my book succeeds. He does not think that TDW-indeterminism is sufficient for L-freedom, and in his comments, he gives an argument for thinking it's not. The argument is a version of the luck objection. I turn to this now.

3.4 Parts of decisions and Kane's version of the luck objection

For the purposes of this section, let's suppose that (a) Ralph made an ordinary torn decision, and (b) the decision was TDW-undetermined. I argued in chapter 3 of my book that, given these assumptions, Ralph's decision was appropriately non-random and L-free. I also argued for the following claim:

(Choice-settles-it) The conscious choice itself *was* the event that settled which option was chosen.

Kane disagrees with this. I think he would admit that if (Choice-settles-it) is true, then we also get the result that the decision was appropriately non-random and L-free. But he denies (Choice-settles-it). More precisely, he thinks that if Ralph's decision was TDW-undetermined, then we do not get the result that the conscious choice settled which option was chosen; rather, we get the result that certain undetermined neural events settled it. Or, pushing this line of thought a bit further, he would presumably endorse the following:

(Quantum-settles-it) Which option was chosen was settled by some quantum events.

I have no problem with (Quantum-settles-it). I just don't think it is incompatible with (Choice-settles-it). Kane seems to think it is. He seems to think that if which option is chosen is settled by some quantum events, then it is not settled by Ralph's decision, or by Ralph. The question, then, is whether (Quantum-settles-it) rules out (Choice-settles-it). I think it does not, and Kane apparently thinks it does. So let me say a few words to motivate my position here.

The first point I want to make is this: It follows from the assumption that Ralph's decision was TDW-undetermined that the relevant quantum events—the ones that settled which option was chosen—were *parts* of the decision. What we do *not* have here is a situation in which prior-to-choice quantum events causally influenced or determined which option was chosen. Rather, given TDW-indeterminism, what we have is a situation in which the conscious choice was *composed of* the relevant quantum events.³ And this is why I think that (Quantum-settles-it) does not rule out (Choice-settles-it).

Kane is not unaware of this—he knows this is my view. He just doesn't seem to think it matters that the decision was composed of the relevant quantum events. I think it does matter, and I want to try to say why. To this end, let me change the example for a moment. Suppose that Johnny and Ernie are having a contest to see who can throw a baseball more accurately. They set up a target with a hole in the middle, and they agree that they will each get ten throws and that hitting the target will count for one point and throwing the ball through the hole will count for two points. Ernie goes first and gets a total of ten points on his ten throws. Johnny goes next, and after nine throws, he has nine points. Finally, on Johnny's last throw, the ball goes through the hole in the target, so that he wins the contest. In this scenario, it seems right to say that the outcome of the contest was settled by Johnny's last throw. But which event in particular settled it? Here are two different views that one might hold:

The baseball view (BV): The outcome of the contest was settled by the event of the baseball going through the hole in the target (we might call this event *the baseball event*).

The simples view (SV): The outcome of the contest was settled by the event (or the several events) of the particles that compose the baseball (or the *simples* that compose the baseball) going through the hole in the target.

Does SV rule out BV? I say it does not. It seems to me that the claim that SV rules out BV leads quickly to the compositional nihilist view that there is simply no such event as the baseball event, i.e., no such event as that of the baseball going through the hole in the target. Compositional nihilism about objects tells us that there are no such things as composite objects like baseballs, so that the only objects that exist in the region of the baseball are some simples arranged baseballwise. Likewise, compositional nihilism about events tells us that there are no such things as composite events like the baseball event, so that the only events that exist in the relevant region of spacetime are events involving metaphysical simples, e.g., the events mentioned in SV. My claim here is that if we reject these radical nihilist views—i.e., if we assume that the baseball and the baseball event exist—then we have no choice but to maintain that SV does not rule out BV. It is simply metaphysical mumbo-jumbo to simultaneously maintain that (i) the baseball and the

³ You might want to say that the quantum events that settled which option was chosen form a *subset* of the events that composed the decision; for you might want to say that the “settling” events were all undetermined events. This will not matter here.

baseball event exist and (ii) BV and SV are competing theories. For if the baseball event exists at all, then it is so intimately related to the events mentioned in SV (in particular, it is *composed* of them) that BV and SV are not genuinely competing claims. In other words, if the baseball event exists, and if SV is true, so that the outcome of the contest was settled by the relevant particles or simples going through the hole in the target, then it is *also* true that the outcome was settled by the *baseball* going through the hole in the target. This much seems obvious to me.

Let's return now to Ralph's decision. The above considerations suggest that if composite events like Ralph's decision exist at all, then (Quantum-settles-it) and (Choice-settles-it) are not in competition. In other words, if Ralph's decision exists, and if it is true that the several quantum events that compose Ralph's decision settle which option is chosen, then it is *also* true that Ralph's decision settles which option is chosen.

Now, again, if the relevant quantum events occurred *before* the decision and *caused* it to go the way that it went, then what I am saying here would not be true. But if Ralph's decision was TDW-undetermined, then this is not the situation. The situation is rather that the decision was *composed* of the quantum events, so that the decision just *was* the relevant bunch of quantum events. But if the decision just *was* the relevant bunch of quantum events, and if the quantum events jointly settled which option was chosen, then the *decision* settled which option was chosen. Again, it seems to me that it would be metaphysical mumbo-jumbo to maintain that (a) the decision existed, and (b) the quantum events that composed the decision jointly settled which option was chosen, and (c) the decision did *not* settle which option was chosen.

One might respond to this argument by endorsing compositional nihilism about events and, hence, denying that Ralph's decision existed at all. But in the present context, I do not think it matters whether nihilism is true. For it seems to me that if nihilism were true, my argument (suitably reworded to allow for the truth of nihilism) would still go through. In this case, there would not be any such things as Ralph or Ralph's decision, but there would still be simples arranged Ralphwise, and there would still be simple events arranged Ralph's-decision-wise, and I think I could just reformulate my argument in terms of these simple events. In short, I could argue that in ordinary human cases, whenever events are arranged TDW-undetermined-torn-decision-wise, they are also arranged appropriately-non-random-L-free-decision-wise. Now, I obviously cannot run through the argument for this here, but (a) I think it is pretty clear that this argument could just parallel the original argument in my book, and (b) in the present context, I don't think I *need* to run through this argument because I think it is reasonable in this context to take the existence of composite events like decisions as a working background assumption.

Let me say one more thing about Kane's version of the luck objection. I don't understand why he thinks this objection applies to my view but not his. Suppose a torn decision satisfies the Kane-conditions. Isn't the outcome of the decision still going to be settled by quantum events? And if it is, then why doesn't the objection that Kane raises against my view apply equally to his own view?

References

- Balaguer, M. (2010). *Free will as an open scientific problem*. Cambridge, MA: MIT Press.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 68, 5–20.
- Hausman, D. (1998). *Causal asymmetries*. Cambridge, UK: Cambridge University Press.
- Kane, R. (2012). Balaguer's *Free will as an open scientific problem*. *Philosophical Studies* (this issue).
- McKenna, M. (2001). Review of John Martin Fischer and Mark Ravizza's *Responsibility and control*. *Journal of Philosophy*, XCVIII(2), 93–100.
- McKenna, M. (2012). The metaphysical importance of the compatibility question: Comments on Mark Balaguer's *Free will as an open scientific problem*. *Philosophical Studies* (this issue). doi: [10.1007/s11098-012-9897-4](https://doi.org/10.1007/s11098-012-9897-4).
- Pereboom, D. (2012). Mark Balaguer's event-causal libertarianism. *Philosophical Studies* (this issue).
- Velleman, J. D. (1992). What happens when someone acts? *Mind*, 101, 462–481.
- Vihvelin, K. (2004). Free will demystified: A dispositional account. *Philosophical Topics*, 32, 427–450.