

Mark Balaguer

Contents

Introduction	232
Is the Determinism Question Relevant to the Free-Will Question?	232
Is the Determinism Question Relevant to the Libertarian Question?	236
Is There Any Good Reason to Doubt that Our Torn Decisions Are TDW-Undetermined?	243
Is There Any Good Reason to Believe Universal Determinism?	243
Is There Any Good Reason to Believe Macro-Level Determinism?	244
Is There Any Good Reason to Believe Neural Determinism?	246
Is There Any Good Reason to Believe Torn-Decision Determinism?	247
Conclusion	249
Cross-References	249
References	249

Abstract

This paper begins with an argument for the claim that the compatibilism question (i.e., the question of whether free will is compatible with determinism) is less relevant than it might seem to questions about the metaphysical nature of human decision-making processes. Next, libertarianism (i.e., the view that human beings possess an indeterministic, libertarian sort of free will) is defended against a number of objections, and it is argued that there's a certain subset of our decisions (which can be called *torn* decisions) for which the following is true: If these decisions are appropriately undetermined at the moment of choice, then they are also free in a libertarian sort of way. This is an extremely important

M. Balaguer

Department of Philosophy, California State University, Los Angeles, CA, USA

e-mail: mbalagu@exchange.calstatela.edu

and surprising result; it entails that the question of whether libertarianism is true reduces to the straightforwardly empirical question of whether our torn decisions are in fact undetermined (in the appropriate way) at the moment of choice. Finally, the paper ends by arguing that as of right now, there is no compelling empirical evidence on either side of this question. In other words, the question of whether our torn decisions are appropriately undetermined is an open empirical question. And from this, it follows that the question of whether libertarianism is true is also an open empirical question.

Introduction

I will do two things in this paper. Ultimately, I will discuss how much evidence there is for various kinds of determinism that might be relevant to the question of whether we humans have free will. Before I do that, however, I will discuss the issue of whether the determinism question is even relevant to questions about the kinds of freedom we have.

Is the Determinism Question Relevant to the Free-Will Question?

Let *determinism* (or as I will sometimes call it, *universal determinism*) be the thesis that every event is causally necessitated by prior events together with causal laws. *Prima facie*, this thesis seems to be incompatible with the thesis that human beings have free will. For if it was already causally determined a billion years ago that my life would take the exact course that it's actually taken – that all of my actions and decisions would turn out exactly as they have – then it would seem to follow that I don't have free will. But some philosophers deny this; they endorse *compatibilism*, i.e., the thesis that free will is compatible with determinism. Compatibilists usually try to motivate their view by (a) providing an analysis of the notion of free will and (b) arguing that, given their analysis, it follows that free will is compatible with determinism. For instance, Hume famously argued that free will is compatible with determinism because free will is essentially just *the ability to do what you want*. Hume put his analysis like this (1748, p. 104):

By liberty, then, we can only mean *a power of acting or not acting according to the determinations of the will*; that is, if we choose to remain at rest, we may; if we choose to move, we also may.

Putting this into contemporary lingo (and altering it somewhat), we arrive at the following definition:

A person *S* is *Hume-free* iff *S* is capable of acting in accordance with his or her choices and of choosing in accordance with his or her desires; i.e., iff it is the case that if he or she chooses to do something, then he or she does it, and if (all things considered) he or she wants to make some choice, then he or she does make that choice.

Hume's argument for compatibilism is based on the following two claims:

- (i) Hume-freedom captures the ordinary notion of free will (that is, Hume-freedom *is* free will).
- (ii) Hume-freedom is compatible with determinism.

The only controversial claim here is (i). In other words, (ii) is entirely obvious. Hume-freedom requires only that our actions flow out of our decisions and our decisions flow out of our desires, but this could be the case even if all of our desires and decisions and actions are causally determined. To appreciate this, it's sufficient to notice that it could be that (a) it was causally determined by the big bang and the laws of nature that we would all have the desires that we actually do have and (b) our desires causally determine our decisions and our decisions causally determine our actions.

Given this, incompatibilists (i.e., those who think that free will is incompatible with determinism) are forced to reject premise (i) of the Humean argument. And that's exactly what they do; they reject the Humean analysis of free will, and they endorse instead a *libertarian* analysis. There are a few different ways to define libertarian-freedom (or as we can also call it, *L-freedom*), but one way to do this is as follows:

A person is *libertarian-free* (or *L-free*) if and only if he or she makes at least some decisions that are such that (a) they are both undetermined and appropriately nonrandom and (b) the indeterminacy is relevant to the appropriate nonrandomness in the sense that it *generates* the nonrandomness, or *procures* it, or *enhances* it, or *increases* it, or something along these lines.

A lot needs to be said about what appropriate nonrandomness consists in, but the basic idea is that in order for a decision to count as appropriately nonrandom, the agent in question has to be centrally involved in the decision. Different philosophers might say slightly different things about what exactly this amounts to, but I think most libertarians would say that the most important requirement for appropriate nonrandomness is that the agent in question has to *author and control* the decision; i.e., it has to be *her* decision, and she has to control which option is chosen. (Other requirements for appropriate nonrandomness might include some kind of *rationality* and what Kane (1996) calls *plural* authorship, control, and rationality; but none of this will matter in what follows.)

In any event, the main point here is that incompatibilists disagree with Humean compatibilists about what free will *is*. Both parties can agree that Hume-freedom is compatible with determinism and L-freedom isn't. But they disagree about whether free will is Hume-freedom or L-freedom.

It's also important to note that the Humean analysis of free will isn't the only compatibilist analysis in the literature. Alternative compatibilist analyses (of not just free will, but moral responsibility) have been put forward by a number of different philosophers, e.g., P.F. Strawson (1962); Frankfurt (1971); Watson (1975); Wolf (1990); Fischer (1994); Wallace (1994); Mele (1995); and McKenna (2012), to name just a few. And, of course, each different analysis of free will gives us a different kind of freedom; thus, e.g., we can define *Frankfurt-freedom*, *Watson-freedom*, *Fischer-freedom*; and so on.

So the question of whether free will is compatible with determinism essentially boils down to the question of whether free will is L-freedom or one of these compatibilist kinds of freedom.¹ Now, this might seem pretty straightforward, but the question is notoriously difficult, and there's no consensus on what the right answer is. There are numerous arguments on both sides of the debate – e.g., there's the Frankfurt-case argument for compatibilism, first articulated in Frankfurt (1969); there's the consequence argument for incompatibilism (the locus classicus of this argument is van Inwagen (1975), but see also Ginet (1966) and Wiggins (1973)); there's the manipulation argument for incompatibilism (see, e.g., Pereboom (2001)) – but none of these arguments have proved really compelling, and at the present time, it seems fair to say that, as a community, we just don't know whether free will is compatible with determinism.

Given this, you might think we should conclude that we don't know the answer to the main question of the present section – i.e., that we don't know whether

The determinism question: Is determinism true?

is relevant to

The do-we-have-free-will question: Do humans have free will?

More specifically, you might think we should say that in order to figure out whether the determinism question is relevant to the do-we-have-free-will question, we first need to figure out what free will is. If free will is L-freedom, then the determinism question *is* relevant to the do-we-have-free-will question, and if free will is some compatibilist kind of freedom, then the determinism question *isn't* relevant to the do-we-have-free-will question.

There's a sense in which this is right, but it seems to me that it's also misleading. For there's another way to think about the issues here, and on this other way of conceptualizing things, less turns on the compatibilism debate than the above remarks suggest. In particular, it seems to me that regardless of whether compatibilism is true, we already know right now that the determinism question is highly relevant to an important question about the nature of human freedom. To appreciate this, notice first that the do-we-have-free-will reduces to (or is subsumed by, or some such thing) the following two more fundamental questions: *The what-is-free-will question:* What is free will? That is, is it Hume-freedom, or L-freedom, or what?

The which-kinds-of-freedom-do-we-have question: Which kinds of freedom do humans have? That is, do they have L-freedom?; and do they have Hume-freedom?; and do they have Frankfurt-freedom?; and so on. (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

¹You might think that for some of the so-called compatibilist kinds of freedom in the literature, it's actually not obvious that they really are compatible with determinism. If this is right, then if it also turned out that one of these kinds of freedom was free will, then we couldn't settle the compatibilism question by merely determining that free will was the given kind of freedom; we would also need to determine whether the given kind of freedom was compatible with determinism.

here might fail to *be* free will, according to the correct answer to the what-is-free-will question.)

We can think of the latter question here (i.e., the which-kinds-of-freedom-do-we-have question) as the *metaphysical* part of the do-we-have-free-will question – i.e., the part that’s actually about the nature of human beings and human decision-making processes. The former question isn’t really about *us* at all; it’s a *semantic* question.² But notice now that the determinism question is obviously relevant to the which-kinds-of-freedom-do-we-have question because it’s relevant to the *libertarian question*, i.e., the question of whether humans are L-free. (Actually, one might deny that the determinism question is relevant to the libertarian question; I’ll discuss this in the next section, but for now, I will ignore this.) In any event, this is what I had in mind when I said that there’s another way to think about the issues here. For if we assume that the determinism question is indeed relevant to the libertarian question, then without even answering the compatibilism question (or the what-is-free-will question), we arrive at the result that the determinism question is relevant to an interesting and important question about the nature of human freedom, namely, the libertarian question, i.e., the question of whether we’re L-free.

Now, I suppose you might claim that if it turns out that L-freedom isn’t free will (i.e., that L-freedom isn’t the referent of the ordinary term “free will”), or if it turns out that L-freedom isn’t required for moral responsibility, then the libertarian question is, in fact, *not* interesting or important. But this is just false. The question of whether we possess an indeterministic, libertarian sort of freedom is *intrinsically* interesting and important. In other words, even if L-freedom isn’t required for moral responsibility, and even if it isn’t the referent of the ordinary term “free will,” the question of whether we actually *possess* L-freedom is itself an interesting and important question about the nature of human beings and human decision-making processes. (Of course, there are *some* people who *aren’t* interested in the question of whether we’re L-free, but so what? – that’s true of every question. There are lots of people who aren’t interested in whether Alpha Centauri has planets, or what the ultimate laws of physics are, or whether humans are morally responsible for their actions. That doesn’t make these questions uninteresting or unimportant.) In any event, since the libertarian question is itself an interesting and important question about the nature of human freedom, it follows that, regardless of whether compatibilism is true, the determinism question is relevant to an interesting and important question about human freedom, namely, the libertarian question, i.e., the question of whether we’re L-free.

²Some people would say that the what-is-free-will question is essentially equivalent to the question, “Which kind(s) of freedom are required for moral responsibility?” But (a) I think it can be argued that the which-kinds-of-freedom-are-required-for-moral-responsibility question is *also* a semantic question (because it’s just a subquestion of the what-is-moral-responsibility question), and (b) even if it’s not a semantic question, it’s pretty clearly not about the metaphysical nature of human decision-making processes.

Is the Determinism Question Relevant to the Libertarian Question?

Libertarianism is often defined as the view that (a) human beings possess L-freedom and (b) L-freedom is free will. But in what follows, I will be concerned with thesis (a) only, and so I will take libertarianism to be the view that humans are L-free, and when I talk about “the libertarian question,” I will have in mind the question of whether we are L-free.

Now, *prima facie*, it seems obvious that the determinism question is relevant to the libertarian question. After all, the libertarian question is just the question of whether we’re L-free, and L-freedom requires indeterminism, and so it seems obvious that the determinism question is relevant to the libertarian question. But you might question this. You might think we can know by logic alone that we’re *not* L-free because the notion of L-freedom is *incoherent*. And if this is right, then the question of whether determinism is true is actually not relevant to the question of whether we’re L-free.

The key claim here is obviously that libertarianism is incoherent. This point has been argued by a number of different philosophers – see, e.g., Hobbes (1651), Hume (1748), and Hobart (1934) – but the reasoning is always very similar. One way to formulate the argument here is as follows:

Anti-libertarian argument: Any event that’s undetermined is uncaused, and so it *just happens* – i.e., it happens randomly. Thus, if we insert an undetermined event into a decision-making process, we’re inserting a *random* event into that process. But given this, it’s hard to see how there could be any undetermined events in our decision-making processes that increase (or generate, or enhance, or whatever) appropriate nonrandomness. Appropriate nonrandomness has to do with the agent being in control. How could this be increased by inserting a random event into the process? It seems that it couldn’t, and so it seems that libertarianism couldn’t be true.

How libertarians respond to this argument is largely determined by the kind of libertarianism they endorse. Broadly speaking, there are three different kinds of libertarian views. First, there are *event-causal* views, which hold that undetermined L-free decisions are *nondeterministically caused* by prior events (I think the best way to fill this view in is to say that these decisions are *probabilistically caused* by prior events, most notably events having to do with the agent’s reasons); event-causal views have been developed by, e.g., van Inwagen (1983), Kane (1996), Ekstrom (2000), and Balaguer (2010). Second, there are *noncausal* libertarian views, which hold that L-free choices are completely uncaused (see, e.g., Ginet (1990)). And third, there are *agent-causal* views, which hold that L-free decisions are caused but not by prior events; rather, they’re directly caused by *persons* via a special causal relation known as *agent causation*; views of this kind have been endorsed by, e.g., Reid (1788), Chisholm (1964), R. Taylor (1966), C.A. Campbell (1967), Thorp (1980), Rowe (1987), O’Connor (2000), and Clarke (1993).

In this section, I will briefly sketch an event-causal response to the above anti-libertarian argument. Most libertarians who have tried to respond to the anti-libertarian argument have done so by trying to explain how our decisions could be simultaneously undetermined and appropriately nonrandom. But I think

libertarians can motivate a much stronger claim than this. I think they can motivate the following thesis:

(L) There's an important subset of our decisions (I'll call them *torn decisions*, and I'll characterize them shortly) for which the following is true: If they're undetermined in the right way, then they're also appropriately nonrandom (i.e., we author and control them), and the indeterminacy in question procures the nonrandomness, and so they're L-free.

This is a really strong result; if it's right, what it shows is that the anti-libertarian argument gets things exactly backwards; more precisely, it shows that a certain kind of indeterminism is sufficient for the truth of libertarianism. And, of course, it also shows that, contrary to what the anti-libertarian argument suggests, the determinism question is definitely relevant to the libertarian question.

Before I argue for (L), I first need to say what a torn decision is, and I need to say what the relevant sort of indeterminacy is, i.e., I need to say exactly how a torn decision needs to be undetermined in order to be L-free. Thus, let me start by saying 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, and whether to drive or bike to the office, and so on. 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're defined in terms of their phenomenology. Thus, we know from experience that we do make torn decisions, and it's an empirical question whether any of these decisions are undetermined in the right way.)

Next, let me define the relevant sort of indeterminacy, i.e., the sort that's 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 this, we can say that *TDW-indeterminism* is the view that some of our torn decisions are TDW-undetermined. And now, given all of these definitions, I can reformulate thesis (L) as follows:

Central-Libertarian-Thesis: 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 nonrandom and L-free. Or more succinctly: *If TDW-indeterminism is true, then libertarianism is true.*

I argued for this thesis at length in a recent book (2010). I can't rehearse the whole argument here, but I'd like to provide brief formulations of two (related) arguments for Central-Libertarian-Thesis.

The first argument is easier to articulate if we assume a weak, token-token mind-brain identity theory – or more precisely, if we assume that ordinary human decisions are neural events. I don't 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 don't matter; option O could be something important like a new job, or it could be something trivial like a chocolate ice cream cone. 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, then 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).
- B. The choice flowed out of Ralph's conscious reasons and thought in a nondeterministically event-causal way.
- C. Nothing external to Ralph's conscious reasons and thought had any significant causal influence (at the moment of choice – i.e., 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 first argument for Central-Libertarian-Thesis is based on the observation that, given (A)–(C), it seems to make sense to say that Ralph authored and controlled his decision. For (A)–(C) seem to give us the twofold result that (i) *Ralph did it and* (ii) *nothing made him do it*; and, intuitively, it seems that (i) and (ii) are sufficient for authorship and control.

Now, to get the result that Ralph's decision is appropriately nonrandom and L-free, we also need to argue that (a) his decision satisfies the other conditions for appropriate nonrandomness, aside from authorship and control (i.e., rationality, the plurality conditions, and so on), and (b) the fact that Ralph's decision is TDW-undetermined *procures* the result that it's appropriately nonrandom and L-free. Point (a) is actually very easy to argue for; I don't have the space to get into this here, but see Balaguer (2010). Point (b), on the other hand, should already be clear from the above argument; for the fact that Ralph's decision was TDW-undetermined played a crucial role in the argument for the claim that he authored and controlled the decision. It's *because* the decision was TDW-undetermined that we get the result that nothing made Ralph choose O over P. Now, it's important to note that the idea here isn't that TDW-indeterminacy actively *generates* authorship and control; the idea is rather that it *blocks a destroyer* of authorship and control. The destroyer of authorship and control would be a moment-of-choice causal influence from something external to the

agent's conscious reasons and thought. But TDW-indeterminacy rules out the possibility of such a destroyer – if a torn decision is TDW-undetermined, then at the moment of choice, nothing external to the agent's conscious reasons and thought comes in and causally influences which option is chosen – and this is why TDW-indeterminacy can be seen as *procuring* authorship and control.

My second argument for Central-Libertarian-Thesis is based on the fact that when we make torn decisions, it *feels* as if we author and control them. The argument can be put like this:

1. The only initially plausible reason to doubt the phenomenology of our torn decisions – i.e., the only reason to doubt our feeling that we author and control these decisions – is that it might be that, unbeknownst to us, our torn decisions are causally influenced (at the moment of choice) by events that are external to our conscious reasons and thought in a way that's inconsistent with the idea that we author and control these decisions. (For example, it could be that our torn decisions are deterministically caused by wholly non-mental brain events that precede our torn decisions in our heads.) But
2. If our torn decisions are TDW-undetermined, then they're *not* causally influenced (at the moment of choice) by anything external to our conscious reasons and thought. Thus,
3. The assumption that our torn decisions are TDW-undetermined seems to eliminate the only initially plausible worry we might have about the accuracy of the phenomenology of our torn decisions. Therefore, it seems plausible to conclude that
4. *If* our torn decisions are TDW-undetermined, then the phenomenology of our torn decisions is accurate and, hence, we author and control these decisions; moreover, it should be clear that the TDW-indeterminacy is *procuring* the authorship and control here, and so we get the result that if our torn decisions are TDW-undetermined, then they're also appropriately nonrandom and L-free.³ In other words, we get the result that
5. Central-Libertarian-Thesis is true.

The two arguments for Central-Libertarian-Thesis that I just articulated are obviously very quick, and there are a number of different worries that one might have about them. I won't be able to respond to all of these worries here, but I'd like to say a few words about two of them. I'll start with this one:

The Rollback Objection: Suppose that Ralph is torn between two options, O and P, and eventually chooses O in a torn decision sort of way. And now suppose that God "rolls back" the universe and "replays" the decision. If the decision is undetermined at the moment of choice, then it seems that the decision might very well go differently the second time around, even if everything about the past – in particular, everything about Ralph and his reasons – remained the same. Indeed, if the decision is TDW-undetermined, then it seems that if God "played" the decision 100 times, we should expect that Ralph would choose

³Actually, to fully motivate (4), we would also need to argue that if our torn decisions are TDW-undetermined, then they satisfy the other requirements for appropriate nonrandomness, i.e., rationality and the plurality conditions. But, again, this point is easy to argue; see Balaguer (2010), Sects. 3.3.4–3.3.5.

O and P about 50 times each. But given this – given that Ralph would choose differently in different “plays” of the decision, without *anything* about his psychology changing – it’s hard to see how we can maintain that Ralph authored and controlled the decision. It seems to be a matter of *chance* or *luck* what he chose, and to the extent that this is right, it seems that Ralph didn’t author or control the choice.

The first point I want to make in response to this objection is that it simply doesn’t follow from the fact that Ralph would choose differently in different “plays” of the decision that he didn’t author or control the decision. There is no inconsistency in claiming that (a) Ralph chooses differently in different plays of the decision and (b) in each of the different plays of the decision, it is *Ralph* who does the choosing and who authors and controls the choice. Indeed, given that Ralph is making a *torn* decision, the hypothesis that it’s *him* who’s making the decision (or who’s authoring and controlling the decision) seems to *predict* that he would choose differently in different plays of the decision. It would seem very suspicious if he always chose the same option in the various plays of the decision; in that scenario, it would be plausible to think: “That can’t be a coincidence; something must be *causing* him to choose that way; and since (by assumption) his conscious reasons and thought aren’t causing this, it must be something else, e.g., a random, non-mental event in his nervous system, or a subconscious mental state.” But if Ralph chose *differently* in different plays of the decision, that would fit perfectly with the hypothesis that the choice is flowing from *him*, or from his conscious reasons and thought; for since Ralph is making a torn decision, we know by assumption that he is neutral between his two live options, at least in his conscious thought. Thus, it seems to me that since Ralph is torn, if he chose differently in different plays of the universe, that wouldn’t undermine the hypothesis that he authors and controls the decision; on the contrary, it would *confirm* that hypothesis.

(You might think that if there was a fixed probability that Ralph was going to choose option O – or, more specifically, if there was a 0.5 probability that he would choose O and a 0.5 probability that he would choose P – then it was just a matter of luck that he in fact *did* choose O, and so it couldn’t be that he authored and controlled the decision. My response to this is that if Ralph’s decision was TDW-undetermined, then (i) it was *his reasons* that caused it to be the case that the probabilities in question were 0.5, and (ii) despite the fact that there were fixed probabilities here, it is still true that the choosing of O over P was done *by Ralph*, because the event in question *was* a Ralph-consciously-choosing event, and this event wasn’t causally influenced by anything external to Ralph’s conscious reasons and thought.)

The second objection I want to consider here is the following:

The Agent-Causal Objection: The notion of control that you’re working with (and the notion of authorship too, but let’s just focus on the case of control) is too weak. Something more substantive is needed for control. In particular, it seems that something like agent causation is needed. In other words, when someone makes a torn decision, in order for it to be the case that the agent in question controls which option is chosen, it needs to be the case that he or she *causes* the given option to be chosen.⁴

⁴Pereboom raises a worry like this about my view in his (forthcoming).

I think it can be argued that agent causation just *isn't* needed for authorship and control, but I won't try to argue this point here. Rather, I want to argue that in the present context, the question of whether agent causation is required for authorship and control doesn't really matter. To bring this out, let me start by distinguishing two different kinds of control – *causal-control* and *noncausal-control* – where the former requires agent causation (or something like it) and the latter doesn't. I won't try to give precise definitions of these two notions; all I'll say (and all that will matter here) is that when I use the term “noncausal-control,” I'm talking about a kind of control that applies to ordinary torn decisions if they're 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 isn't 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. Beyond this (and beyond the fact that causal-control requires agent causation and noncausal-control doesn't), it won't matter how exactly these two kinds of control are defined. But for the sake of argument, let's pretend that we've got two precisely defined kinds of control here. Given this, one question we might ask is the following:

The what-is-control question: What is control? That is, which of the various kinds of control that we find in the literature is *real* control? Is causal-control real control? Is noncausal-control? Are both? Is neither?

But why should libertarians care about this question? They don't need to claim that if our torn decisions are TDW-undetermined, then they're 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 they need is this:

(*) If our torn decisions are TDW-undetermined, then they're authored and controlled by us and appropriately nonrandom 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.”

Now, don't take me to be saying more than I am here. I'm not saying that libertarians can define authorship and control and L-freedom *however they want to*; they can't just define these terms in ridiculously weak ways and then claim victory. I don't need to argue that the kind of L-freedom I've articulated – the kind that we get if our torn decisions are TDW-undetermined (i.e., the kind that involves noncausal-control) – 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 above arguments for Central-Libertarian-Thesis do motivate (*). Let's return to Ralph's decision. If it's TDW-undetermined, then (a) the 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. This might not give us every kind of L-freedom you might have wanted, but it clearly gives us *one important kind* of L-freedom – a kind that libertarians can hang their hats on and that's 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's the event with a me-consciously-choosing-now phenomenology.

There is obviously a lot more to say about all of this. In Balaguer (2010), I develop the arguments for Central-Libertarian-Thesis a lot more thoroughly, and I respond to a number of different objections to these arguments. For instance, I say more by way of response to the luck objection; and I respond to the worry that the kind of L-freedom I've been describing here isn't robust enough to give us moral responsibility; and I also respond to the worry that this kind of L-freedom isn't worth wanting because torn decisions are trivial. Unfortunately, though, I don't have the space to discuss these issues here.

In any event, if what I've argued in this section is correct, then we have the result that if TDW-indeterminism is true (i.e., if some of our torn decisions are TDW-undetermined), then libertarianism is also true (i.e., humans are L-free). But given this, it follows pretty quickly that the question of whether libertarianism is true just reduces to the question of whether TDW-indeterminism is true.⁵ And if this is right, then we have an answer to the main question of this section: The determinism question is definitely relevant to the libertarian question.

⁵To argue for this, libertarians need to argue that if TDW-indeterminism *isn't* true, then libertarianism isn't true either – i.e., that if our torn decisions aren't TDW-undetermined, then we aren't L-free. Now, you might doubt this, because you might think that even if our torn decisions aren't L-free, some of our non-torn decisions could be L-free. But it's pretty easy to argue – and I do argue for this point in Balaguer (2010) – that if none of our torn decisions is L-free, then it's very likely that we don't make any L-free choices at all. The more important worry about the above thesis (i.e., the thesis that if TDW-indeterminism isn't true, then libertarianism isn't true either) is that even if our torn decisions aren't TDW-undetermined (i.e., even if they aren't *wholly* undetermined), they could still be *partially* undetermined in a certain way. To say that a torn decision is partially undetermined in the sense I have in mind here – or what comes to the same thing, partially determined – is to say that, at the moment of choice, factors external to the agent's conscious reasons and thought causally influence (but don't causally determine) which tied-for-best option is chosen. I think it can be argued that if our torn decisions are partially undetermined in this way, then they're also *partially* L-free. Thus, to be precise, what we need to say here is not that if TDW-indeterminism isn't true, then we aren't L-free, but that if TDW-indeterminism isn't true, then we aren't *fully* L-free. And so to get the result that if TDW-indeterminism isn't true, then libertarianism isn't true, we need to be clear that we're defining libertarianism as the view that humans are *fully* L-free.

Is There Any Good Reason to Doubt that Our Torn Decisions Are TDW-Undetermined?

We found in the last section that the question of whether libertarianism is true boils down to the question of whether TDW-indeterminism is true (i.e., the question of whether some of our torn decisions are TDW-undetermined). In this section, I want to discuss the question of whether we have any good reason to reject TDW-indeterminism. I think it's pretty obvious that, at present, we don't have any good reason to *endorse* TDW-indeterminism, but you might think we have good reason to *reject* it, because you might think we have good reason to endorse some deterministic thesis that's incompatible with TDW-indeterminism. I will argue in this section, however, that as of right now, we don't have any good reason to believe any such deterministic thesis.

Is There Any Good Reason to Believe Universal Determinism?

Let *universal determinism* (or UD) be the thesis that I've been calling "determinism" – i.e., the thesis that every event is causally necessitated by prior events together with causal laws. It's pretty easy to see that as of right now, we don't have any good reason to believe UD. For UD is true only if all quantum events are determined, and as of right now, we don't have any good reason to believe that all quantum events are determined. The first point to be made here is that quantum mechanics (or QM) contains probabilistic laws; it tells us, for instance, that if an electron is spin-up in a particular direction x , then it's in a superposition state with respect to its spin in the orthogonal direction y , and if we measure it for spin in the y -direction, then it will collapse into either a spin-up state or a spin-down state, and there's a 0.5 probability that it will collapse into a spin-up state and a 0.5 probability that it will collapse into a spin-down state.

Now, of course, the fact that QM contains probabilistic laws of this kind does not by itself give us reason to doubt UD; for it could be that there are hidden facts (or as physicists say, *hidden variables*) about electrons that are spin-up in the x -direction that determine whether they will collapse into a spin-up state or a spin-down state when measured for spin in the y -direction. But the problem is that there is no good evidence for the existence of hidden variables of this kind, and so, for all we know, it could just as easily be that when electrons that are spin-up in the x -direction are measured for spin in the y -direction, *nothing* determines whether they collapse into a spin-up state or a spin-down state; in other words, for all we know, it could be that events of this kind – i.e., events involving quantum wave-function collapses – are genuinely undetermined.

This is not to say that we have good reason to endorse an indeterministic view of these events. Rather, it's to say that as of right now, we have no good reason to reject an indeterministic view. In other words, the question of whether these quantum collapse events are genuinely undetermined or just apparently undetermined (i.e., really determined) is an open question. There is simply no good evidence on either side of the debate. Or to put the point differently, there's

no good evidence for any (deterministic or indeterministic) interpretation of QM. An interpretation of QM is essentially a theory of what's going on in quantum collapse events of the above kind; there are numerous interpretations in the literature, some deterministic and some indeterministic, but at present, there isn't any evidence for any of them, and more generally, there isn't any compelling reason to endorse a deterministic or an indeterministic view of quantum collapse events.⁶

Is There Any Good Reason to Believe Macro-Level Determinism?

If the arguments of the previous section are cogent, then there's no good reason to believe universal determinism. But in order to undermine TDW-indeterminism and libertarianism, you don't need to motivate universal determinism. Since TDW-indeterminism is about torn decisions only, you could undermine this thesis by arguing for the much weaker claim that all torn decisions are determined. One way to do this would be to point out that torn decisions are macro-level events and then argue for *macro-level determinism*, i.e., the view that all macro-level events are determined. Or, alternatively, you could undermine TDW-indeterminism by arguing for what might be called *virtual macro-level determinism*, i.e., the view that while it may be that some macro-level events are strictly undetermined (because they're composed of micro-level events, some of which are undetermined), it's also true that all macro-level events are, if not determined, at least virtually determined (where an event is *virtually determined* iff prior circumstances together with causal laws made it overwhelmingly likely that the given event would occur). In other words, the idea here is that while there may be some micro-level indeterminacies, these all "cancel out" before we get to the macro level, presumably because macro-level phenomena are composed of such large numbers of micro-level phenomena. (It should be clear that virtual macro-level determinism would undermine TDW-indeterminism; after all, it entails that all torn decisions are virtually determined, i.e., that for any torn decision, there's a unique option X such that prior events made it overwhelmingly likely that X would be chosen.)

The question I want to ask now is whether we have any good reason to believe macro-level determinism or virtual macro-level determinism. People sometimes claim that there's a good inductive argument for macro-level determinism (see, e.g., Honderich (2002)⁷). We might put the argument here as follows:

1. All the macro-level events that we have encountered have been causally determined by prior events together with causal laws. Therefore,
2. Macro-level determinism is true – i.e., all macro-level events are determined.

⁶Of course, there are people who favor certain interpretations over others, but there is pretty widespread agreement among those who work on the foundations of quantum mechanics that we do not have any solid evidence for any of the various interpretations and that when people embrace these interpretations, they are engaged in speculation.

⁷Actually, Honderich thinks we can use arguments like the one in the text to motivate not just macro-level determinism but universal determinism as well.

But this argument is misguided. In particular, premise (1) is unmotivated, controversial, and question begging. We encounter all sorts of macro-level events that, for all we know, could be undetermined – coin tosses, events in which a person contracts chicken pox from someone else, events in which macro-level measuring devices reveal quantum wave-function collapses, human decisions, chimp decisions, parakeet decisions, temper tantrums, etc. Now, of course, determinists have a story to tell about how it *could be* that events like these are deterministic; e.g., they can claim that if, say, Jack and Jill were both exposed to chicken pox and only Jack fell ill, this would not undermine determinism because it could be that there were hidden physical variables at work in the situation (e.g., factors having to do with the physical well-being of Jack and Jill, or the duration of their exposures, or whatever) that determined that Jack would contract the disease and Jill would not. And likewise for events of the other kinds listed above; determinists can say that events like coin tosses and decisions could be determined even if they don't seem determined to us, because it could be that there are hidden determining factors at work in such cases. I agree; for all we know, it *could be* that events of the above kinds are determined. But in the present context, this is entirely irrelevant. What advocates of the argument in (1)–(2) need isn't a story about how it *could be* that events of the above kinds are determined; what they need is a positive argument for the claim that, in fact, such events *are* determined.

But I take it that determinists don't have an argument of this kind. The argument they used to give here is that any apparently indeterministic behavior in macro-level systems must really be deterministic, because such systems are made up of micro-level systems whose behavior is deterministic. But this argument is no good, because (as we've seen) we currently have no more reason to believe micro-level determinism than macro-level determinism.

Now, I suppose one might respond here by claiming that every time we go looking for deterministic explanations, we find them. But this is just false. It's not just that we don't currently have deterministic explanations of events of the above kinds; it's that we haven't the foggiest idea how to proceed in trying to construct and justify such explanations.

The situation with virtual macro-level determinism is similar. One might try to argue for this view by saying something like the following:

- 1'. All the macro-level events that we've encountered have been either determined or virtually determined. Therefore,
- 2'. Virtual macro-level determinism is true – i.e., all macro-level events are either determined or virtually determined.

But this argument is flawed in the same way the (1)–(2) argument is flawed. In short, the problem is that (1') is unmotivated, controversial, and question begging. There are lots of macro-level events – coin tosses, quantum-measurement events, decisions, and so on – that, for all we know, are neither determined nor virtually determined. In order for virtual macro-level determinists to motivate an inductive argument of the above kind, they would need to provide positive reasons for thinking that events like coin tosses and decisions and quantum measurements

are, in fact, either determined or virtually determined. But at present, there is simply no good reason to believe this.

Finally, it's worth noting here that if the remarks in this section are correct, they suggest not just that the above inductive arguments are noncogent, but that, right now, we don't have any good reason to believe macro-level determinism or virtual macro-level determinism.

Is There Any Good Reason to Believe Neural Determinism?

Since torn decisions are presumably neural events, you might think that we could undermine TDW-indeterminism (and hence libertarianism) by uncovering reasons to believe *neural determinism* (i.e., the view that all neural events are determined) or *virtual neural determinism* (i.e., the view that all neural events are either determined or virtually determined in the sense defined above). But, in fact, we don't have any good reason to believe either of these theses. If current neuroscientific theory were deterministic (or virtually deterministic), then we might be able to motivate neural determinism (or virtual neural determinism). But current neuroscientific theory is *not* deterministic (or virtually deterministic). Indeed, it treats a number of different neural processes probabilistically – e.g., synaptic transmission and spike firing. Consider, e.g., the following passages from a recent textbook on neuroscience (Dayan and Abbott 2001):

- I. . . [synaptic] transmitter release is a stochastic process. Release of transmitter at a presynaptic terminal does not necessarily occur every time an action potential arrives and, conversely, spontaneous release can occur even in the absence of the depolarization due to an action potential. (p. 179)
- II. Because the sequence of action potentials generated by a given stimulus varies from trial to trial, neuronal responses are typically treated statistically or probabilistically. For example, they may be characterized by firing rates, rather than as specific spike sequences. (p. 9)

It is worth noting that some aspects of the indeterminacies in these processes (or the apparent indeterminacies, as the case may be) are caused by the indeterminacy (or apparent indeterminacy) in another process, namely, the opening and closing of ion channels. Now, to be sure, by treating these processes probabilistically, neuroscientists don't commit themselves to the thesis that, in the end, they are genuinely indeterministic. But the important point here is that they aren't committed to determinism either. The question of whether these processes are genuinely indeterministic simply isn't answered by neuroscientific theory. Indeed, it is a standard view among those who work in this area that for at least some of these processes (e.g., the opening and closing of ion channels), this isn't even a neuroscientific question, because it is already clear right now that there could not be deterministic neuroscientific explanations of the phenomena. In other words, the idea is that (a) from the point of view of neuroscience, these processes might as well be undetermined but (b) it *could* be that there are underlying

deterministic *physical* explanations of the phenomena. Thus, the question of whether there actually are such explanations is not a neuroscientific question at all; it is rather a question of physics, because the issue comes down to questions about the behavior of the elementary physical particles involved in the neural processes.

It sum, then, it seems to me that neuroscientific theory is neither deterministic nor virtually deterministic, and so it doesn't give us any reason to believe neural determinism or virtual neural determinism. And given this, it seems safe to conclude that as of right now, we don't have any good reason to believe neural determinism or virtual neural determinism.

Is There Any Good Reason to Believe Torn-Decision Determinism?

Finally, you might try to undermine TDW-indeterminism by arguing for *torn-decision determinism* (i.e., the view that all torn decisions are determined) or *virtual torn-decision determinism* (i.e., the view that all torn decisions are either determined or virtually determined). Or, of course, you could try to argue directly against TDW-indeterminism; i.e., you could try to give a direct argument for the claim that none of our torn decisions is TDW-undetermined. In this section, I will respond to one such argument, an argument based on the work of Benjamin Libet.

(It's worth noting that the argument based on Libet's work isn't the only argument against TDW-indeterminism that one might consider here. Another important argument – we might take it to be an argument for something like virtual torn-decision determinism – is based on Max Tegmark's (2000) argument for the claim that if there are any neural superposition states, they couldn't survive long enough to be affected by neural processes. One might also construct arguments against TDW-indeterminism by using considerations having to do with situationism (see, e.g., Isen and Levin 1972, Milgram 1969, and Nelkin 2005), the sluggishness of consciousness (see, e.g., Velmans 1991 and Wegner 2002), or the way in which humans are often out of touch with the real underlying reasons for their actions (see, e.g., Festinger 1957). In Balaguer (2010), I argue that none of these considerations provides us with any good reason to reject TDW-indeterminism; but unfortunately, I don't have the space to pursue any of this here.)

In any event, let's consider the argument against TDW-indeterminism that's based on Libet's work. Libet's studies were a follow-up to a neuroscientific discovery from the 1960s, in particular, the discovery that voluntary decisions are preceded in the brain by a slow change in electrical potential known as the *readiness potential* (see, e.g., Kornhuber and Deecke 1965). Libet's studies were an attempt to establish a timeline for the readiness potential, the conscious intention to act, and the act itself (see, e.g., Libet et al. 1983, and Libet 2002). His results suggest that the readiness potential appears about 350–400 milliseconds before the conscious intention to act and about 550 milliseconds before the act itself.

And given this, one might argue against TDW-indeterminism in something like the following way:

1. Conscious decisions are preceded by nonconscious brain processes (namely, the readiness potential) and are, in fact, nonconsciously initiated. Therefore, it seems likely that
2. Torn decisions are at least causally influenced by prior-to-choice nonconscious brain processes, and so they are not TDW-undetermined; indeed, they might be determined, or virtually determined, by prior-to-conscious-choice brain processes.

In other words, the idea here is that our torn decisions couldn't be TDW-undetermined because (to borrow Henrik Walter's (2001) phrasing) the "neural machinery" for starting our decisions is already up and running before our conscious thinking enters the picture.

One might try to attack the argument in (1)–(2) by questioning (1), but I won't pursue this strategy here. What I want to argue instead is that even if (1) is true, it does not give us any good reason to accept (2). The first point to note here is that we don't know what the *function* of the readiness potential is. In particular, it would be an unmotivated assumption to suppose that, in torn decisions, the readiness potential is part of a causal process that's relevant to which option is chosen. There are plenty of other things the readiness potential could be doing, aside from this. One way to appreciate this is to notice that libertarianism is perfectly consistent with the idea that various things involved with our torn decisions might be causally determined. In particular, a torn decision could be L-free even if it was determined in advance that (i) a torn decision would occur, and (ii) the choice would come from among the agent's reasons-based tied-for-best options, and (iii) the moment-of-choice probabilities of these options being chosen were all roughly even. The only thing that needs to be undetermined, in order for a torn decision to be L-free, is *which tied-for-best option is chosen*. Given this, here are two stories libertarians could tell about what the readiness potential could be doing (there are other stories as well – see, e.g., Mele (2009) – but these two will do):

Model A: (a) The readiness potential is part of the causal process leading to the *occurrences* of torn decisions, and this has nothing whatsoever to do with which option is chosen; and (b) which option is chosen is in fact TDW-undetermined. (A similar point, though a bit different, has been made by Haggard and Eimer – see, e.g., their (1999) as well as Haggard's contribution to Haggard and Libet (2001).)

Model B: (a) The readiness potential is part of the process whereby our reasons cause our decisions, and (b) in connection with torn decisions, this process doesn't determine which option is chosen; rather, it deterministically causes it to be the case that the choice will come from among the reasons-based tied-for-best options (and perhaps also that the moment-of-choice probabilities of these options being chosen are all roughly even).

Now, models A and B are both highly controversial, and as of right now, I don't think we have any good reason to endorse either of them. But the important

point here is that as of right now, we don't have any good reason to reject them either; in particular, the available evidence concerning the readiness potential doesn't give us any good reason to reject them. More generally – and in the present context, this is the really important point – as of right now, there is no reason to think that, in torn decisions, the readiness potential is part of a causal process that's relevant to the issue of which tied-for-best option is chosen. There is simply no evidence for this, and so the existence of the readiness potential doesn't give us any good reason to suppose that, in torn decisions, which tied-for-best option is chosen is causally influenced by prior-to-choice nonconscious brain processes.

Conclusion

In the last section, I responded to a variety of arguments against TDW-indeterminism. There are, of course, other arguments that one might attempt here, but I don't think any of them are cogent. In other words, at the present time, I don't think we have any good reason to reject TDW-indeterminism. And as I pointed out above, I don't think we have any good reason to *endorse* TDW-indeterminism either. Thus, if this is right, then the question of whether TDW-indeterminism is true is a wide open question. But earlier I argued that the question of whether libertarianism is true (i.e., the question of whether humans are L-free) reduces to the question of whether TDW-indeterminism is true. Thus, it seems that as of right now, the libertarian question is an open question. And in particular, it's an open *empirical* question. For (a) the question of whether we're L-free turns on the question of whether TDW-indeterminism is true, and (b) TDW-indeterminism is a straightforward empirical hypothesis about the causal histories of our torn decisions.

Cross-References

- ▶ [Consciousness and Agency](#)
- ▶ [Free Will and Experimental Philosophy: An Intervention](#)
- ▶ [Mental Causation](#)
- ▶ [No Excuses: Performance Mistakes in Morality](#)

References

- Balaguer, M. (2010). *Free will as an open scientific problem*. Cambridge, MA: MIT Press.
- Campbell, C. A. (1967). *In defense of free will*. London: Allen & Unwin.
- Chisholm, R. (1964). Human freedom and the self. Reprinted in Watson (1982), pp. 24–35.
- Clarke, R. (1993). Toward a credible agent-causal account of free will. *Nous*, 27, 191–203.

- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.
- Ekstrom, L. W. (2000). *Free will: A philosophical study*. Boulder: Westview Press.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Palo Alto: Stanford University Press.
- Fischer, J. M. (1994). *The metaphysics of free will: A study of control*. Oxford: Blackwell.
- Frankfurt, H. (1969). Alternate possibilities and moral responsibility. *Journal of Philosophy*, 66, 829–839.
- Frankfurt, H. (1971). Freedom of the will and the concept of a person. *Journal of Philosophy*, 68, 5–20.
- Ginet, C. (1966). Might we have no choice? In Lehrer (Ed.), (pp. 87–104).
- Ginet, C. (1990). *On action*. Cambridge: Cambridge University Press.
- Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126, 128–133.
- Haggard, P., & Libet, B. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies*, 8, 47–63.
- Hobart, R. E. (1934). Free will as involving determinism and inconceivable without It. *Mind*, 43, 1–27.
- Hobbes, T. (1651). *Leviathan*. New York: Collier Books. Reprinted 1962.
- Honderich, T. (Ed.). (1973). *Essays on freedom of action*. London: Routledge and Kegan Paul.
- Honderich, T. (2002). Determinism as true, compatibilism and incompatibilism as false, and the real problem. In Kane (Ed.), (pp. 461–476).
- Hume, D. (1748). *An inquiry concerning human understanding*. Indianapolis: Bobbs-Merrill. Reprinted 1955.
- Isen, A., & Levin, P. (1972). Effect of feeling good on helping. *Journal of Personality and Social Psychology*, 21, 384–388.
- Kane, R. (1996). *The significance of free will*. New York: Oxford University Press.
- Kane, R. (2002). *The oxford handbook of free will*. New York: Oxford University Press.
- Kornhuber, H., & Deecke, L. (1965). Hirnpotentialänderungen bei willkurbewegungen und passiven bewegungen des menschen. *Pfluegers Arch Gesamte Physiologie Menschen Tiere*, 284, 1–17.
- Libet, B. (2002). Do we have free will? In Kane (Ed.), (pp. 551–564).
- Libet, B., Gleason, C., Wright, E., & Pearl, D. (1983). Time of conscious intention to Act in relation to cerebral potential. *Brain*, 106, 623–642.
- McKenna, M. (2012). *Conversation and responsibility*. Oxford: Oxford University Press.
- Mele, A. (1995). *Autonomous agents*. New York: Oxford University Press.
- Mele, A. (2009). *Effective intentions*. New York: Oxford University Press.
- Milgram, S. (1969). *Obedience to authority*. New York: Harper and Row.
- Nelkin, D. (2005). Freedom, responsibility, and the challenge of situationism. *Midwest Studies in Philosophy*, XXIX, 181–206.
- O'Connor, T. (2000). *Persons and causes*. New York: Oxford University Press.
- Pereboom, D. (2001). *Living without free will*. Cambridge: Cambridge University Press.
- Pereboom, D. (forthcoming). Mark Balaguer's event-causal libertarianism. *Philosophical Studies*.
- Reid, T. (1788). *Essays on the active powers of the human mind*. Cambridge, MA: MIT Press. Reprinted 1969.
- Rowe, W. (1987). Two concepts of freedom. *Proceedings of the American Philosophical Association*, 62, 43–64.
- Strawson, P. F. (1962). Freedom and resentment. *Proceedings of the British Academy*, 48, 1–25.
- Taylor, R. (1966). *Action and purpose*. Englewood Cliffs: Prentice-Hall.
- Tegmark, M. (2000). The importance of quantum decoherence in brain processes. *Physical review E*, 61, 4194.
- Thorp, J. (1980). *Free will*. London: Routledge & Kegan Paul.
- van Inwagen, P. (1975). The incompatibility of free will and determinism. Reprinted in Watson (1982), pp. 46–58.
- van Inwagen, P. (1983). *An essay on free will*. Oxford: Clarendon.

- Velmans, M. (1991). Is human information processing conscious? *The Behavioral and Brain Sciences*, 14, 651–669.
- Wallace, R. J. (1994). *Responsibility and the moral sentiments*. Cambridge, MA: Harvard University Press.
- Walter, H. (2001). *Neurophilosophy of free will*. Cambridge, MA: MIT Press.
- Watson, G. (1975). Free agency. *Journal of Philosophy*, 72, 205–220.
- Watson, G. (1982). *Free will*. Oxford: Oxford University Press.
- Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
- Wiggins, D. (1973). Towards a reasonable libertarianism, In Honderich (Ed.), (pp. 31–61).
- Wolf, S. (1990). *Freedom within reason*. New York: Oxford University Press.