Synthese (2009) 168:1–21
DOI 10.1007/s11229-009-9459-2

Why there are no good arguments for any interesting version of determinism

Mark Balaguer

Received: 28 September 2007 / Accepted: 13 January 2009 / Published online: 12 February 2009

Abstract This paper considers the empirical evidence that we currently have for various kinds of determinism that might be relevant to the thesis that human beings possess libertarian free will. Libertarianism requires a very strong version of indeterminism, so it can be refuted not just by universal determinism, but by some much weaker theses as well. However, it is argued that at present, we have no good reason to believe even these weak deterministic views and, hence, no good reason—at least from this quarter—to doubt that we are libertarian free. In particular, the paper responds to various arguments for neural and psychological determinism, arguments based on the work of people like Honderich, Tegmark, Libet, Velmans, Wegner, and Festinger.

Keywords Determinism · Libertarianism · Free will · Neural determinism · Psychological determinism

1 Introduction

In this paper, I will consider the empirical evidence that we currently have for various kinds of determinism that might be relevant to the thesis that human beings possess libertarian free will. As we will see, libertarianism requires a very strong version of indeterminism, so it can be refuted not just by universal determinism, but by some much weaker theses, e.g., claims to the effect that all of our conscious decisions are significantly causally influenced (not necessarily genuinely determined) by prior events in certain kinds of freedom-damaging ways. I will argue that, at present, we have no good reason to believe even these weak deterministic views and, hence, that

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

we have no good reason—at least from this quarter—to doubt that we are libertarian free.

We can say (somewhat loosely) that a person has libertarian free will, or *L-freedom*, if and only if some of her decisions are such that (i) they are undetermined, (ii) they are appropriately non-random, and (iii) the indeterminacy is relevant to the non-randomness in the sense that it generates it, or procures it, or some such thing. The reason this is loose is that more needs to be said about the kinds of indeterminacy and non-randomness that are required here. I will have a good deal to say about indeterminacy below, but since the notion of appropriate non-randomness will not be very relevant to this paper, there is no need to say much about it; suffice it to say that it has to do mainly with *agent involvedness* (e.g., one might put this in terms of the person in question *controlling* which option is chosen, or being the *author* of the decision, or some combination of claims like this).

It's important to note that the kind of L-freedom I've got in mind is entirely naturalistic. Thus, e.g., talk of persons or agents controlling decisions should not be understood as involving any sort of dualism or irreducible agent causation. The idea here is that we naturalistic humans are L-free just in case (a) some of the neural events that are our decisions are appropriately undetermined, and (b) the indeterminacy gives rise to an appropriate sort of non-randomness. (This isn't to say that L-freedom is naturalistic or materialistic by *definition*; one could also develop a dualistic version of libertarianism; but I think the best versions of libertarianism are materialistic, and in any event, the kind of L-freedom I've got in mind here is materialistic.)

Finally, note too that libertarianism—which I'm taking here to be the thesis that humans are L-free[1]—does not involve the claim that *all* human decisions are undetermined. Libertarians can allow that many of our decisions are (or may be) determined, and in particular, they can maintain that in cases where a person's reasons clearly pick out a unique best option, it is perfectly fine (i.e., consistent with libertarianism) if the choice of the given option is causally determined by the person's having the reasons in question. What libertarianism requires (or at any rate, what the sort of libertarianism that I've got in mind requires) is that some of our *torn decisions* are L-free—i.e., satisfy conditions (i)–(iii) above—where a torn decision is a decision in which the person in question (a) has reasons for two or more options and feels torn as to which set of reasons is strongest, i.e., has no conscious belief as to which option is best, given her reasons; and (b) decides without resolving this conflict and, thus, has the experience of "just choosing" from a list of reasons-based tied-for-best options. (Sometimes torn decisions are very important—e.g., one might have a deadline by which to accept or decline an exciting job offer in an undesirable city, and one might have to decide while feeling utterly torn—but it is important to note that we make more mundane torn decisions all the time; e.g., we often decide what to order for dinner while feeling torn between two or more options.)[2,3,4]

[1] Some would define libertarianism as the view that (a) humans are L-free and (b) free will is L-freedom. But that's not how I'm using the term here.

It is widely believed that libertarianism could not be right, because (a) it requires a freedom-enhancing sort of indeterminism, and (b) it seems that to introduce an undetermined event into a decision-making process would be to introduce an element of *randomness*, which, if anything, would be freedom-damaging, not freedom-enhancing. Several libertarians (e.g., Kane 1996; Ginet 1990) have tried to argue that indeterminacy and appropriate non-randomness are in fact compatible, so that libertarianism is at least coherent. I have argued (2004, forthcoming) for the much stronger claim that the right kind of indeterminacy actually *entails* the right kind of non-randomness, so that the question of whether humans are L-free just reduces to the question of whether some of our torn decisions are undetermined in the right way. Even if I'm wrong about this, however, the question of whether any of our torn decisions are appropriately undetermined is still an important question, for it is still relevant to the question of whether libertarianism is true, even if it isn't the whole question, as it is on my view. At any rate, this is the question that I want to investigate in this paper—whether any of our torn decisions are undetermined in the right way.

In Sect. 2, I will describe the kind of indeterminism that's required for L-freedom (I'll call it *TDW-indeterminism*), and I will articulate a number of different deterministic theses (or pseudo-deterministic theses, as the case may be, since one might doubt that these theses should be thought of as versions of determinism) that are inconsistent with TDW-indeterminism. Then in Sects. 3–5 I will argue that, at present, we have no good reason for believing any of these kinds of determinism and, more generally, no good reason to reject TDW-indeterminism. And if we combine this with the plausible assumption that we also have no good reason to *endorse* TDW-indeterminism, we obtain the result that the central question here—whether our decision-making-processes are undetermined in the way that's required for L-freedom—is simply an open empirical question.

Before going on, I want to make three general points. First, I'm going to argue that, at present, we have no "good arguments" for any of the kinds of determinism I'll be discussing. If you interpret expressions like 'good argument' and 'good reason'

[2] I don't mean to suggest that torn decisions are the only kinds of decisions that can be L-free. On the libertarian view I've got in mind, the primary instances of L-freedom involve torn decisions, but libertarians can argue that if some of our torn decisions are L-free, then other kinds of decisions can be L-free as well. I can't go into this here, but see Kane (1996) and my (forthcoming) for more on this issue.

[3] Torn decisions might seem similar to what Kane has called self-forming actions (SFAs), but there are a few important differences. For instance, SFAs are by definition undetermined and L-free, and so it's an open question whether there actually exist any such things. Torn decisions, on the other hand, are not defined as being undetermined or L-free; they are defined in terms of their phenomenology; thus, it's clear that there do exist torn decisions—we know this by experience—but it's an open question whether any of our torn decisions are undetermined or L-free.

[4] I'm writing here as if the only possible cases are (a) those in which the agent believes that one of her options is clearly best and (b) those in which the agent is torn between multiple options that seem to her tied for best. But there is actually a continuum of possible cases here. The cases in (a) and (b) are limiting cases, at opposite ends of the spectrum; in between are cases in which the agent is "leaning" toward one or more of her options but still thinks she has other live options. For the sake of simplicity, I will ignore this complication here and concentrate on torn decisions.

weakly, you might take this claim to be stronger than I intend it to be. By a good argument, I mean a *compelling* argument, or an argument that, all things considered, gives us a rational reason to believe the conclusion, or something along these lines. Thus, my central claim here is that, as of right now, all things considered, we ought (rationally) to remain neutral on the question of whether any of the relevant kinds of determinism are true. But in saying this I do *not* mean to suggest that there are no considerations that provide any coherent rationale—or any arguments that ought to be taken seriously—for any of these kinds of determinism. Second, since this paper provides a sort of survey—in particular, since it covers a variety of different arguments for several different deterministic theses—I will not be able to go into as much depth on any of the various issues as I would like. This is a bit unfortunate, but it's just an unavoidable constraint on survey pieces like this. Finally, one might be thinking that the conclusion of this paper is somewhat uninteresting because it's already widely accepted that we don't have any good reason to believe determinism. But I am not just arguing that we have no good reason to believe determinism; indeed, I won't be saying much of anything about universal determinism; what I'm going to argue is that we have no good reason to believe any of the "pseudo-deterministic" theses that would rule out TDW-indeterminism and libertarian free will; and this, I think, is not something that people already believe.

2 TDW-indeterminism and its competitors

I now want to describe the sort of indeterminacy that's needed for L-freedom. In particular, I want to make three points here. First, I'm going to assume that what's needed for L-freedom is indeterminacy *at the moment of choice*, not prior to choice. I think it can be argued that this is what's needed, but I won't argue this here because, in the present context, the point is relatively unimportant. The arguments of this paper are couched in terms of the issue of moment-of-choice indeterminacy, but almost everything I say here could be reformulated in a way that would make it relevant to the issue of prior-to-choice indeterminacies as well.[5]

The second point I want to make here is that libertarianism is consistent with the thesis that the mere *occurrences* of our torn decisions are determined (and also with the thesis that when a torn decision is made, it is causally determined that the person in question will choose from a specific list of options, namely, her reasons-based tied-for-best options). What must be undetermined in order for a torn decision to be L-free is precisely *which option is chosen*. Thus, e.g., if you're buying a soda and you make a torn decision to buy coke rather than root beer, then it's OK (i.e., consistent with the thesis that the decision was L-free) if it was determined (by your reasons) that you were going to make a decision, and that you were going to choose between coke and root beer, so long as it wasn't determined that you were going to choose coke.

[5] When I speak of "the moment of choice", I do not mean to suggest that I think choices occur instantaneously; on the contrary, I think our decisions are very likely temporally extended events, i.e., events that are spread out in time.

The third and final point I want to make here is related to the fact that, conceptually speaking, there can be different degrees of causal determination. In any torn decision, the agent feels neutral between her tied-for-best options, given all of her conscious reasons and thought. Thus, we might describe this by saying that the *reasons-based probabilities* of the various live options being chosen are all even, or at least roughly even. But it could be that factors external to the agent's conscious reasons and thought (e.g., unconscious compulsions, or wholly non-mental brain events that precede the decision in the agent's head) causally influence the choice and wholly or partially determine which option is chosen. Indeed, there is a continuum of conceptually possible cases here. At one end of the spectrum, which option is chosen is wholly undetermined; that is, the moment-of-choice probabilities of the various reasons-based tied-for-best options being chosen match the reasons-based 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. At the other end of the spectrum, which option is chosen is causally determined by prior events together with causal laws. And in between, there is a continuum of possible cases where the moment-of-choice probabilities of the various tied-for-best options being chosen are altered, away from the reasons-based probabilities, to a greater or lesser extent, due to some sort of causal input; in connection with these in-between cases, we can say that which option is chosen is partially determined—or, equivalently, partially undetermined. Now, what I argue in my (2004) and my (forthcoming) is that if any of our torn decisions are *wholly* undetermined in the above way, then they're also (wholly) appropriately non-random and L-free; thus, assuming this is right, the question we ought to be concerned with, vis a vis libertarianism, is whether any of our torn decisions are wholly undetermined in this way. (Or perhaps we should just say that this is the *primary* question we should be concerned with; for there is also a question about whether any of our torn decisions are partially undetermined, and this might end up being important, for one might argue that even if our torn decisions aren't wholly undetermined, as long as they're partially undetermined, they will still be partially appropriately non-random and L-free.)

Given all this, we can say that what's required for the sort of libertarianism I'm describing—i.e., for it to be the case that some of our torn decisions are (wholly) L-free—is that the following empirical hypothesis be true:

> *TDW-indeterminism*: Some of our torn decisions are wholly undetermined at the moment of choice, where to say that a torn decision is wholly undetermined at the moment of choice is to say that the moment-of-choice probabilities of the various reasons-based tied-for-best options being chosen match the reasons-based 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.

In this paper, I want to motivate the thesis that we do not currently have any good arguments for or against TDW-indeterminism. Now, I will not argue every point that needs to be argued here. What I'll actually do is block a handful of arguments that one

might use to try to refute TDW-indeterminism. But since I think these are the most promising arguments out there, this will motivate the claim that we do not currently have any good argument against TDW-indeterminism. And if we combine this with the plausible assumption that there are no good arguments in *favor* of TDW-indeterminism, it follows that there are no good arguments for or against that view.[6]

There are a few different strategies one might use to try to argue against TDW-indeterminism. First, one could try to argue for *universal determinism*, the view that all events are causally determined by prior events together with physical laws. Second, one could try to argue for *macro-level determinism*, the view that whatever we say about micro-level events, all macro-level events are determined. (This, of course, would suffice, because our torn decisions are neural events, and so they're macro-level events.) Third, one could try to argue for what we might call *virtual macro-level determinism*, the view that while it may be true 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 virtually determined—i.e., that in every macro-level situation (or experiment), there is a specific possible outcome or event such that prior circumstances together with causal laws make it overwhelmingly likely that the given event will occur. In other words, the idea here is that while there may be some micro-level indeterminacies, these all "cancel out", or "disappear", before we get to the macro level, presumably because macro-level phenomena are composed of such large numbers of micro-level phenomena. (Clearly, virtual macro-level determinism would undermine TDW-indeterminism, because it entails that all torn decisions are virtually determined—i.e., that for any torn decision, there is a unique option X such that it is overwhelmingly likely that X will be chosen.) Fourth and fifth, one could zero in even more, not just to the macro level, but all the way to the neural level; that is, one could try to argue for *neural determinism*, the view that all neural events are determined, or *virtual neural determinism*, the view that all neural events are virtually determined in the above sense. Sixth and seventh, one could narrow one's attention even further and try to argue for *torn-decision determinism* or *virtual torn-decision determinism*. And finally, eighth, one could try to give a direct argument against TDW-indeterminism itself.

In this paper, I will (for the most part) ignore universal determinism and concentrate on blocking arguments for the various other kinds of determinism and virtual determinism defined above. I do this because it's pretty widely accepted that we don't right now have any good arguments for universal determinism. The reasoning here seems to be that (a) an argument for universal determinism would have to proceed by establishing that all micro-level events are determined,[7] and (b) we have no good arguments for micro-level determinism, because quantum mechanics contains probabilistic laws,

[6] I think there are parallel arguments suggesting that we have no good reason to believe or disbelieve the thesis that some of our torn decisions are *partially* undetermined, and given what I say here it will be pretty obvious how these arguments go, but I will not explicitly spell them out.

[7] Of course, one might reject (a) and try to argue for universal determinism by arguing that our experience with macro-level events suggests that *all* events are determined. But I will respond to arguments of this sort in Sect. 2.

and there are no good arguments for any deterministic interpretation of these laws.[8] And it's worth noting here that most of the well-known proponents of deterministic interpretations of quantum mechanics would actually grant point (b). For instance, Bohm would not claim that there is any evidence for his interpretation; all he would say—and he's certainly right about this—is that there is no evidence against it.[9]

In Sect. 3, I argue that there are no good arguments for macro-level determinism or virtual macro-level determinism; I concentrate here on an inductive argument put forward most notably by Honderich. In Sect. 4, I argue that there are no good arguments for neural determinism or virtual neural determinism; I concentrate here on issues in neuroscience. And in Sect. 5, I consider what I think are the most plausible strategies one might use to mount an argument in favor torn-decision determinism or virtual torn-decision determinism or against TDW-indeterminism in particular, and I argue that none of the arguments here is good. In particular, I consider (i) an argument based on the work of the physicist Max Tegmark (Tegmark's argument is directed against the Penrose-Hameroff quantum-computer view of the brain, but it's been taken by others to raise a problem for the sort of indeterminism required for free will); (ii) an argument based on the work of Benjamin Libet; and (iii) a cluster of arguments based on the many psychological studies showing that much of our behavior is significantly causally influenced by factors that we have no conscious access to. I have in mind here studies related to situationism, confabulation, and the slowness of consciousness; I do not deny that these studies show that nonconscious factors are more important to the causation of our behavior than we might like to admit; my claim will simply be that these studies don't give us good reason to reject TDW-indeterminism.

In arguing that there are no good arguments for any of the above versions of determinism or virtual determinism, I do not mean to suggest that I think there are good arguments *against* these views. In fact, I don't, and in a few places, I say a bit about this, though I don't need this result here, because the existence of such an argument wouldn't entail an answer to the question of whether TDW-indeterminism is true anyway (although an argument against virtual torn-decision determinism would certainly be a step in the direction of an argument for TDW-indeterminism).

3 Arguments for macro-level determinism or virtual macro-level determinism?

People sometimes try to argue for macro-level determinism by means of an inductive generalization. We might put the argument here in the following way:

(1) All of the macro-level events that we have encountered in our lives 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.

[8] It's commonly claimed that we have good reasons for endorsing an *in*deterministic interpretation of quantum mechanics; I think this is pretty clearly false; the fact of the matter is that there are no good arguments on either side of this debate, because there is no good evidence for any theory of what's really going on in quantum wave-function collapses. I will return to issues related to this in Sect. 5.1.

[9] See Bohm (1952) and Bohm and Hiley (1993). It is sometimes said that deterministic interpretations of quantum mechanics are ruled out by Bell's theorem (1964). This is false. Indeed, Bell himself was a Bohm-style determinist; see, e.g., his (1982).

One might also try to motivate universal determinism in the same way. Ted Honderich tries to do this in the following passage (2002, p. 462):

> In my life so far I have never known a single event to lack an explanation in the fundamental sense, and no doubt your life has been the same. No spoon has mysteriously levitated at breakfast. There has been no evidence at all, let alone proof, of there being no explanation to be found of a particular event. On the contrary, despite the fact that we do not seek out or arrive at the full explanations in question, my experience and yours pretty well consists of events that we take to have such explanations. If we put aside choices or decisions and the like—the events in dispute in the present discussion of determinism and freedom—my life and yours consists of nothing but events that we take to have fundamental explanations. Thus, to my mind, no general proposition of interest has greater inductive and empirical support than that all events whatever, including choices or decisions and the like, have explanations.

In the present section, I will argue that this sort of reasoning is misguided. I will concentrate on the argument in (1)–(2), arguing that our experiences don't even motivate macro-level determinism (and, of course, it follows from this that they don't motivate universal determinism either). Moreover, after undermining the argument in (1)–(2), I will undermine an analogous inductive argument for virtual macro-level determinism.

The problem with the (1)–(2) argument for macro-level determinism is very simple: premise (1) is unmotivated, wildly controversial, and question begging. We encounter all sorts of macro-level events that seem as though they might be undetermined—or more accurately and importantly, that are such that we have no idea whether they are determined or not—e.g., coin tosses, the appearances of unsightly facial blemishes on the mornings of proms, 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, fallings in love, cars running out of gas after being run for 314 miles rather than 314.0001 miles, 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 very easily 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 might be deterministic, even if they seem random to us, because it might be that there are hidden determining factors at work in such cases. I agree; for all we know, it *might* be that events of the above kinds are fully determined. But in the present context, this is entirely irrelevant. What determinists need, in order for the argument in (1)–(2) to have any force, is not 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.

I take it that determinists have no response to this, i.e., that they do not have an argument of the required sort. The argument they used to give is that any apparently

indeterministic behavior of 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 we currently have no more reason to believe micro-level determinism than macro-level determinism.

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 these phenomena; it's that we haven't the foggiest idea how to proceed in trying to construct and justify such explanations.

What about virtual macro-level determinism? Well, there is an inductive argument for this view that's exactly analogous to the (1)–(2) argument. We might put it like this:

(1′) All of the macro-level events that we have encountered in our lives 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—e.g., coin tosses, quantum-measurement events, decisions, and so on—that, for all we know, might be neither determined nor virtually determined. In order for virtual macro-level determinists to motivate an inductive argument of this sort, 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 Honderich-style inductive arguments are noncogent, but that, right now, we have no good reason to believe macro-level determinism or virtual macro-level determinism.

4 Arguments for neural determinism or virtual neural determinism?

In this section, I will refute an argument for neural determinism, and in the process, I will argue that, at present, there are no good arguments for neural determinism or virtual neural determinism, because current neuroscientific theory (and the existing empirical evidence) are perfectly consistent with the falsity of these theories, i.e., with the thesis that some neural events are neither determined nor virtually determined.

Honderich (1988, Chap. 5) claims to have an argument for neural determinism. In particular, he argues for two theses that he thinks motivate neural determinism. The two theses are (a) that "Mental and neural events are intimately connected—each specific type of mental event somehow necessarily occurs with a simultaneous specific type or types of neural event" (1988, p. 269), and (b) that "Neural events are the effects of standard causal sequences," i.e., of "prior neural or other bodily events" (1988, p. 288).

But in fact, these theses don't motivate neural determinism. Thesis (a) is clearly consistent with neural indeterminism. Honderich thinks that thesis (a) creates a problem

for dualistic views of the mind, and he seems to think that indeterministic views of the mind have to be dualistic. But, of course, this is false. There is no reason for neural indeterminists or libertarians to endorse dualism, and indeed, the version of libertarianism described above involves a commitment to a token-token mind-brain identity theory, and it's perfectly consistent with type-type identity theories that entail thesis (a). So, again, thesis (a) poses no threat to the kind of indeterministic view that we're concerned with here.

Similar remarks apply to thesis (b). Most of what Honderich says here does little more than motivate the idea that there's nothing "spooky" about brain processes—in particular, that the brain is not causally influenced by weird, "non-standard", or immaterial forces. Again, this is something that neural indeterminists can and should accept. They can say that neural causation is just like other kinds of physical causation. The only difference between neural determinists and neural indeterminists is that the latter think that some neural events aren't causally determined by prior events. But this doesn't mean they think there is something unusual going on in the brain, because they can of course say that there are lots of non-neural events that aren't causally determined either—e.g., quantum events, quantum-measurement events, unsightly-facial-blemish events, and so on.[10]

Now, there is one reading of thesis (b) according to which it is incompatible with neural indeterminism; in particular, if 'standard causal sequence' means 'deterministic causal sequence', then obviously, thesis (b) is little more than a rewording of neural determinism. And, indeed, it's obvious that this is what Honderich has in mind. But the problem is that he doesn't have any *argument* for thesis (b) on this reading. His discussion motivates nothing stronger than the claim that the brain is a physical, causal system. There isn't a shred of evidence given for the claim that all of the causation involved in the brain is deterministic causation. The only thing he says on this topic is that neuroscientists understand causation as deterministic causation. He writes (1988, p. 292):

> there is no doubt that neuroscience everywhere uses the intuitive and standard conception of a causal circumstance and effect, such that the former necessitates the latter.

But this claim is just straightforwardly false. Or at any rate, the claim that neuroscience treats all neural processes deterministically is straightforwardly false; what the standard neuroscientist's conception of causation is, I don't know, and as far as I can see, it doesn't matter. Current neuroscientific theory treats a number of different neural processes probabilistically, and any decent textbook on neuroscience will point this out. For instance, synaptic transmission and spike firing are both treated probabilistically. One textbook (Dayan and Abbott 2001) puts these points as follows:

[10] One might think that the conjunction of (a) and (b) motivates neural determinism, or virtual neural determinism, even if neither (a) nor (b) does so in isolation. But this is false. As Honderich understands these theses, the conjunction essentially says that (i) some sort of mind-brain materialism is true, and (ii) brain events are caused in ordinary ways. But neural indeterminists can endorse this conjunction and then simply claim that some neural events are probabilistically caused.

> (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 both of these processes are caused by the indeterminacy inherent in another process, namely, the opening and closing of ion channels, which are essentially little gates that let charged ions in and out of cells. Now, to be sure, by treating these processes probabilistically, neuroscientists do not 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.

That this is a standard view among those who work in this area is an empirical claim. My reason for believing it arises out of private correspondences I have had on this point with a few neuroscientists. Dayan, for instance, says that "people would argue that there are good thermal reasons to think that [the opening and closing of ion channels] is truly random. Thus, short of philosophical debates about hidden variables for all forms of randomness in physics, this is some fundamental randomness for which people nearly have evidence." And Sebastian Seung, a neuroscientist at MIT, says that "The question of whether [synaptic transmission and spike firing] are 'truly random' processes in the brain isn't really a neuroscience question. It's more of a physics question, having to do with statistical mechanics and quantum mechanics." And finally, Christof Koch, a Cal Tech neuroscientist, says: "At this point, we do not know to what extent the random, i.e., stochastic, neuronal processes we observe are due to quantum fluctuations (a la Heisenberg) that are magnified by chemical and biological mechanisms or to what extent they just depend on classical physics (i.e., thermodynamics) and statistical fluctuations in the underlying molecules."[11]

It seems, then, that standard neuroscience is *consistent with* neural determinism but that it doesn't come close to entailing (or, indeed, of providing any reason at all to endorse) neural determinism or even virtual neural determinism. In short, the claim

[11] All of the quotes in this paragraph are from private correspondences.

that current neuroscientific theory gives us reason to believe neural determinism or virtual neural determinism is just straightforwardly false.

So, returning to Honderich, I agree with his claim that the brain is an ordinary physical system involving ordinary causal processes. What I deny is that it follows from this that these processes are deterministic. As far as I can see, Honderich just hasn't given any evidence at all for the claim that all brain processes—which, granted, are ordinary, physical, causal processes—are deterministic. And it should be noted that the point here goes beyond a refutation of Honderich's argument; given what we've seen here about current neuroscientific theory, it seems safe to conclude that as of right now, there is no good empirical reason to believe neural determinism or virtual neural determinism.

5 Arguments for torn-decision determinism, or for virtual torn-decision determinism, or against TDW-indeterminism?

Let us now focus our attention all the way down to the level of torn decisions. Even if there are no good arguments for universal determinism, macro-level determinism, virtual macro-level determinism, neural determinism, or virtual neural determinism, one might still think there are good reasons for believing torn-decision determinism or virtual torn-decision determinism—or, more specifically, for rejecting TDW-indeterminism, the thesis that's actually needed for libertarianism to be true. In this section, I will consider what I think are the strongest arguments for these claims, and I will argue that none of them is cogent. In particular, in Sect. 5.1, I respond to an argument that arises out of the work of Max Tegmark; in Sect. 5.2, I respond to an argument that arises out the work of Benjamin Libet; and in Sect. 5.3, I respond to a collection of arguments that arise out of various psychological studies suggesting that much of our behavior is determined by nonconscious phenomena.

5.1 The argument from Tegmark's work

Tegmark (2000) has given an argument that some people have taken as providing strong reason for doubting the existence of any indeterminacies that are relevant to mental events like torn decisions. Tegmark advertises his argument as a refutation of the Penrose-Hameroff view that the brain acts as a quantum computer and that this is centrally important to consciousness (see Penrose and Hameroff 1995 and Hameroff and Penrose 1996); but the argument is really an argument against a specific thesis that's inherent in the Penrose-Hameroff theory, namely, the thesis that there are neural superposition states in the brain that undergo wave-function collapses due to neural processes.

To make sense of this, we need a bit of background on quantum mechanics. The first point to be made here is that quantum mechanics contains probabilistic laws; it tells us, for instance, that if an electron is spin-up in the x-direction, then it is in a *superposition* state with respect to its spin in the orthogonal direction y—where this is a kind of indeterminate state, a sort of "hovering between" spin-up and spin-down—and if we measure the electron for spin in the y-direction, then there is a 0.5 probability

that it will be spin-up and a 0.5 probability that it will be spin-down. It's important to realize that these superposition states cannot be understood epistemologically, in terms of observer ignorance. It's not that the electron is already spin-up or spin-down in the y-direction (prior to our measurement of its spin in that direction) and that we just don't know which it is; the electron isn't spin-up or spin-down in the y-direction prior to measurement. There is a very strong argument for this due to Kochen and Specker (1967).[12] Thus, on the standard view, the superposition state is physically real, and the measurement process involves a *collapse*—what's known as a collapse of the wave function—into a spin-up or a spin-down state. Now, as is well known, there are deterministic and indeterministic interpretations of quantum wave-function collapses, but we do not right now have any good evidence for any of the theories here. The fact of the matter is that, as of right now, we don't know what superposition states really *are*, and we don't know what's really going on in wave-function collapses.[13]

The Penrose-Hameroff theory involves two central ideas. The first, proposed by Hameroff, is that tubulin proteins residing within neurons could, in theory, take on superposition states. Tubulin proteins make up the walls of microtubules, which are tiny hollow cylinders that (along with other structures) make up the "skeletons", or "internal scaffoldings", of neurons. Hameroff's idea is that since tubulin proteins can take on two different shapes, namely, *extended* and *contracted*, they could in principle be in extended-contracted superposition states. The second idea in the Penrose-Hameroff model is that tubulin-protein superpositions are inherently unstable and, therefore, subject to wave-function collapses—or objective reductions, or self-collapses—brought on by a quantum-gravity mechanism suggested by Penrose. Tegmark's argument is aimed not at the specific idea that there are tubulin-protein superpositions that self-collapse due to a Penrose-style quantum-gravity mechanism, but rather at the more general claim that there are neural superpositions, or macro-level brain-state superpositions, that undergo wave-function collapses due to neural processes. The argument, in a nutshell, is that because the brain is so warm and wet, neural superpositions could not survive long enough to be affected by neural processes. More specifically, Tegmark argues that (a) because of the brain's temperature, and because of disturbances within the brain caused by things like ions and water molecules, any neural superpositions that might obtain within the brain would decohere due to the destructive influence of environmental "noise" within about 10^{-13} s; but (b) neurons function, at the fastest, on a time scale of about 10^{-3} s; and so (c) even if there are neural superpositions, they could not avoid decoherence for a long enough period of time to undergo self-collapse.

This argument has been taken to show that there couldn't be any quantum indeterminacies that matter to mental processes like decision making. For instance, Hodgson (2002, p. 107) says that Tegmark's argument, if cogent, shows that "in systems as massive, hot, and wet as neurons of the brain, any quantum entanglement and

[12] The argument is a geometrical proof of the claim that if the probabilistic laws of quantum mechanics are correct (and there is overwhelming empirical evidence for them), then no electron could possibly have a determinate value of spin for every direction in space at the same time.

[13] Of course, there are people who favor certain interpretations over others, but I think 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.

indeterminacies would be eliminated within times far shorter than those necessary for conscious experiences". If this is right, then Tegmark's argument provides some motivation for virtual torn-decision determinism. For (a) as of right now, it seems plausible to suppose that the best hope for torn-decision indeterminism is the existence of quantum indeterminacies in the brain, and (b) on the above reading of Tegmark's argument, it suggests that even if there are quantum indeterminacies in the brain, they disappear quickly and aren't relevant to mental events like torn decisions.

But I want to argue that even if Tegmark's argument is cogent—even if it succeeds in refuting the Penrose-Hameroff view[14]—it *doesn't* show that there are no quantum indeterminacies in the brain that are relevant to mental events like torn decisions. What his argument shows is that there couldn't be any neural indeterminacies of a certain, specific kind, namely, indeterminacies based in macro-level brain-state superpositions that undergo wave-function collapses due to neural processes. But Tegmark's argument does not show that there couldn't be any neural indeterminacies that involve *micro*-level superpositions and wave-function collapses. More specifically, for all Tegmark has argued, it could be that the following model is true:

> *The neural-dependence-on-micro-indeterminacy model of neural/mental indeterminacy*: (i) There are some neural events that are such that some of their constituent micro-level events are undetermined; and (ii) for some of these neural events, some of their macro-level features (e.g., whether there's a synaptic transmission or a neural firing) depend on the outcomes of the constituent indeterministic micro-level events in such a way that prior to the occurrence of the neural event in question, there was a significant probability of its being different, so that the neural event is not only not determined, but also not virtually determined; and (iii) in some cases of this sort, if the undetermined neural event(s) go one way, then there is one sort of mental event (or set of mental events), and if they go another way, then there is a different sort of mental event (or set of mental events). Finally, it could be that in some cases of this sort, (a) if the undetermined micro-level events go one way, we get one neural event (call it N_1), whereas if they go another way, we get a different neural event (call it N_2); and (b) N_1 and N_2 are both torn decisions, but in N_1 the agent in question chooses one of her reasons-based tied-for-best options, whereas in N_2 she chooses another.

[14] One might also try to argue (though I will not do this here) that Tegmark's argument is in fact not cogent. Indeed, Hagan et al. (2002) have tried to do just this. They attack both of the crucial numbers in Tegmark's argument—i.e., the 10^{-13} figure he comes up with for maximum decoherence time and the 10^{-3} figure he comes up with for the minimum time that the superpositions would have to survive to be relevant to neural processes. In connection with the 10^{-13} figure, they argue that Tegmark ignores a variety of factors at work in the vicinity of microtubules that could increase decoherence time, and taking the various considerations that they cite into account, they recalculate the decoherence time and come up with a range of 10^{-5} to 10^{-4} s. In connection with the 10^{-3} figure, on the other hand, they claim that sequences of tubulin-superposition collapses in the range of 10^{-7} to 10^{-6} s could result in neural events in the $10^{-2} - 10^{0}$ range. Thus, in sum, the idea is that the relevant collapses could take about 10^{-7} s (faster than the maximum decoherence time of 10^{-4} s) and that these could give rise to neural events that fall within the slower-than-10^{-3}-s range. I am not qualified to assess the arguments on either side of this dispute, but I think that most of those who are would side with Tegmark.

Clearly, this model could be true even if there are no such things as macro-level superpositions or macro-level wave-function collapses (and even if the brain is not a quantum computer). Thus, it could be true even if Tegmark's argument is entirely cogent. And so Tegmark's argument should not be thought of as establishing the claim that there couldn't be any quantum indeterminacies that are relevant to mental events like torn decisions. It is rather an argument for the thesis that there couldn't be any neural indeterminacies of a certain, specific kind, namely, indeterminacies involving macro-level brain-state superpositions that undergo wave-function collapses due to neural processes.[15]

The neural-dependence-on-micro-indeterminacy model gives us a kind of neural indeterminacy that's consistent with Tegmark's argument but still relevant to mental events like decisions. But it's important to note that I am not claiming that we have good reason to endorse this model. All I'm saying is that as of right now, we have no good reason to doubt it. At any rate, Tegmark's argument doesn't give us a reason to doubt it, and as far as I know, we don't have any good reason at all. (Note, too, that my stance here is made more plausible by the arguments of Sect. 4, which show that current neuroscientific theory and the existing empirical evidence are perfectly consistent with the existence of neural events that are neither determined nor virtually determined.)

It's also important to note how *weak* the neural-dependence-on-micro-indeterminacy model is. In formulating this model, I have tried to capture the most general version of the idea that there are neural/mental indeterminacies that are based in micro-level indeterminacies; that is, I have tried not to give any details about how the view might be filled in, and so I have located what is, I think, the weakest version of the view. And again, my claim is simply that, as of right now, there is no good reason to disbelieve this view.

Having said how my stance here is weak, though, I should also point out something that libertarians who endorsed this stance would be committed to. In order for a libertarianism of the sort sketched in Sect. 1 to be true, it needs to be the case that some of our torn decisions are undetermined *at the moment of choice*. So not just any scenario that involves neural dependence on micro-level indeterminacy will do here. In particular, if the relevant undetermined micro-level event(s) are *prior to*, and not part of, the torn decision, we will not procure the required sort of indeterminacy. The undetermined micro-level events need to be *part of* the torn decision, so that the decision itself (in particular, which option is chosen) is undetermined. Now, this is not to say that the undetermined micro-level events need to be simultaneous with the *whole* of the decision; as long as they were parts of the decision, it would be OK (i.e., consistent with the above kind of libertarianism) if the decision continued for a period of time after the completion of the undetermined micro-level events. But this issue raises an important point: there is no guarantee that in every case there will be a clearly right answer to the question of whether the relevant undetermined micro-level events were (a) prior to the neural/mental event that is the torn decision or (b) part of it.

[15] It's worth noting that whatever others have said, Tegmark himself seems to think that his argument is consistent with the thesis that there are indeterminacies in the brain that survive to the macro level (see, e.g., his 2000, p. 11).

There could be cases where it would clearly be best to say that the relevant micro-level events were prior to the relevant neural/mental event, and there could be cases where it would clearly be best to say that the former were part of the latter; but there could also be cases where there was no clearly right answer to the question, because the cut-off between what was prior to the decision and what was part of it was vague.

In any event, it seems to me that for the reasons given above, Tegmark's results do not give us any good reason to think that any of our torn decisions are determined or virtually determined. Indeed, his results seem to be perfectly consistent with the libertarian thesis that some of our torn decision are wholly undetermined in the manner of TDW-indeterminism.

5.2 The argument from Libet's work

Libet has produced some results that have been taken to generate a problem for free will. In particular, his work seems to generate a problem for the sort of indeterminism that's required for L-freedom. The story here begins with a neuroscientific discovery that goes back to the 1960s. It was discovered then that voluntary decisions are preceded in the brain (by as much as a second) by a slow change in electrical potential; this electrical shift, which is recordable on the scalp, is known as the *readiness potential* (see Kornhuber and Deecke 1965 and Deecke et al. 1976). Building on this work, Libet has tried to establish an exact timeline for the readiness potential, the conscious intention to act, and the act itself (see Libet et al. 1983 and Libet 2002). His results suggest that the readiness potential appears about 350–400 ms before the conscious, experienced intention to act, and about 550 ms before the act itself.

These data have been taken to raise a serious problem for free will. For instance, Henrik Walter (2001, p. 249) writes: "Libet's...findings immediately evoked the query. If we are not aware of our intentions until after the neural machinery for starting the act has already been warmed up, is free will just an illusion?" And Max Velmans (1991, p. 658) says that Libet's work "suggests that conscious volition may be one output from the (prior) cerebral processes that actually select a given response." More specifically, Libet's results seems to generate a problem for libertarian views that are based on TDW-indeterminism. For TDW-indeterminism seems inconsistent with the idea that our torn decisions are determined prior to the moment of conscious volition, and, to use Walter's way of putting things, Libet's findings seem to suggest that the "neural machinery for starting an act" is already up and running before the agent's conscious thought enters the picture.

To make things a bit more precise, one might argue from Libet's findings to the rejection of 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; and
(2) These nonconscious brain processes are presumably *not parts of* the conscious decisions in question. Therefore, it seems likely that
(3) Torn decisions are at least causally influenced by prior-to-choice nonconscious brain processes, and so they are not wholly undetermined in the manner of TDW-indeterminism (and they might even be determined, or virtually determined, by prior-to-conscious-choice brain processes).

One might try to attack this argument by questioning (1) or (2), but I will not pursue this strategy here. What I want to argue instead is that even if (1) and (2) are true, they do not give us any good reason to accept (3).

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 recall (from Sect. 2) 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, (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 count as L-free, is *which option is chosen*. Given this, here are two stories libertarians could tell about what the readiness potential might be doing (there are other stories as well—see, e.g., Mele (2006)—but these two will do):

Model A: It might be that (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 wholly undetermined in the manner of TDW-indeterminism. (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: It might be that (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, of course, the views contained in models A and B are 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 seem to 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 reason to suppose that, in torn decisions, which option is chosen is causally influenced, prior to the moment of conscious choice, by nonconscious processes.

5.3 Arguments from psychology

There are a number of different psychological studies that generate doubts about human free will. For instance, there are studies that suggest that:

1. Consciousness is sluggish—i.e., conscious awareness of various actions and processes lags behind the processes themselves (see, e.g., Velmans 1991 and Wegner 2002).
2. People are often mistaken about why they perform various actions and, indeed, confabulate reasons for their actions. (There is a mountain of evidence for this; see, e.g., Festinger 1957.)
3. Our actions and behaviors are often significantly influenced by situational factors that, intuitively, seem relatively unimportant (see, e.g., Isen and Levin 1972 and Milgram 1969, and for a discussion, see Nelkin 2005).
4. Conscious choices can be causally influenced by magnetic stimulation to the brain (Brasil-Neto 1992).

I think these studies are important and that they reveal all sorts of interesting and often depressing facts about humans. But I don't think they give us any good reason to reject TDW-indeterminism. In general, the problem is that these studies either (a) don't tell us anything at all about torn decisions or (b) don't tell us anything universalizable about them. Let me say a bit about this.

In connection with point 1, the evidence concerns processes for which, intuitively, it's not surprising that awareness lags behind action—things like the processing of incoming speech and knee-jerk reactions in emergency situations (e.g., the jerking of a steering wheel to one side to avoid hitting a child who has suddenly run in front of your car). Whatever we end up saying about cases like these, there is no evidence right now that in torn decisions, consciousness lags behind the actual selection of one of the reasons-based tied-for-best options.[16]

In connection with 2 and 3, there are a number points to make. First, while there is overwhelming evidence for the thesis that much of our behavior is causally influenced by factors to which we have no conscious access, most of the studies here are irrelevant to our question, because they concern actions and behavior that aren't prefaced by conscious decisions (and certainly aren't prefaced by *torn* decisions). This is a point that is too infrequently noticed in discussions of what these studies tell us about free will. The fact of the matter is that very few of our actions are prefaced by decisions. (And this, by the way, is good; imagine what a nightmare it would be to have to consciously will all your actions. Your inner monologue during a stroll might go something like this: move the left foot forward; ok, now the right; easy going—not too far with that foot—land right there, just before the hole, and so on.) Libertarianism and TDW-indeterminism are concerned with cases involving conscious decisions, in particular, torn decisions, i.e., cases in which we feel torn, pause for at least a moment, and then consciously and intentionally choose one of our tied-for-best options. Very

[16] In making this claim, I do not mean to suggest that people are not morally responsible or L-free in cases like these. In other words, there is no suggestion here that the only time people are L-free is when they make torn decisions. See note 2 for more on this.

few of the experiments related to points 2 and 3 have any bearing at all on these cases. On the other hand, I think it's fair to say that the sum total of all the experiments here provides ample evidence for thinking that at least sometimes—and perhaps very often—our torn decisions are causally influenced by nonconscious factors. I don't think anyone who knows the psychological literature would want to deny this. But the point I want to make here is that libertarians don't *need* to deny it. All they need to do is endorse TDW-indeterminism, which says only that *some* of our torn decisions are not inappropriately caused by nonconscious factors. (It might be better to say that libertarians need to maintain that a *significant percentage* of our torn decisions are TDW-undetermined—or something along these lines; but it seems plausible to assume that there is some degree of regularity here, so that if some of our torn decisions are TDW-undetermined, then many of them are.) In any event, this point is important, because as interesting and important as the various psychological experiments are, they don't come close to motivating the conclusion that *all* of our torn decisions are causally influenced by nonconscious factors in ways that are inconsistent with TDW-indeterminism. Think of a typical day: you might make torn decisions about whether to have eggs or cereal for breakfast, whether to exercise before going to work, whether to go to your office or work at home, whether to take the freeway or surface streets to work, whether to meet a friend for lunch or eat in your office to get more work done, whether to work late or go to a movie, and so on. Does the psychological literature really suggest that *none* of the decisions that we make of this kind is TDW-undetermined? The answer, I think, is that it does nothing of the sort. The evidence that we presently have seems perfectly consistent with the thesis that a significant percentage of our torn decisions are TDW-undetermined.

Finally, the Brasil-Neto study mentioned in point 4 might seem particularly relevant to TDW-indeterminism, because it's directly concerned with something like torn decisions (subjects were told to raise either their left or right finger when they heard a click, and unbeknownst to them, their choices were correlated with whether the left or right sides of their brains were magnetically stimulated)[17]. But in fact, this study doesn't undermine TDW-indeterminism at all. The study is a sort of real-life version of the various alien-manipulation thought experiments that philosophers often discuss. But even if aliens can manipulate our torn decisions, it doesn't follow that when aliens aren't present, our torn decisions aren't libertarian free.

In sum, then, I don't think any of the arguments or considerations discussed in this paper give us any good reason to reject TDW-indeterminism. More generally, I don't know of any good arguments for or against TDW-indeterminism. If this is right, then the question of whether TDW-indeterminism is true is a wide open empirical question.

Now, even if TDW-indeterminism is true, one might still doubt that human beings are L-free. Indeed, as I pointed out in Sect. 1, it is widely believed that even if our decisions were undetermined, they wouldn't be L-free, because the insertion of undetermined events into our decision-making processes couldn't generate or increase freedom, because undetermined events are essentially random, and randomness is antithetical to freedom. But, again, I have argued (2004, forthcoming) that this

[17] One might call decisions like this *Buriden's-ass decisions*, rather than torn decisions; in the latter, the agent has *reasons* for both options, and more importantly, the reasons are *different*.

traditional line of thought gets things exactly backwards. I argue not just that indeterminacy *could* increase freedom, but that the right kind of indeterminacy—i.e., TDW-indeterminacy—*would* increase freedom. More specifically, I argue that if our torn decisions are TDW-undetermined then they're also L-free, so that the question of whether we're L-free just reduces to the question of whether TDW-indeterminism is true. And if we combine this with the arguments of the present paper, we obtain the result that the question of whether human beings are L-free is a wide open empirical question.

Acknowledgements I would like to thank Robert Kane, Timothy O'Connor, Carl Ginet, Dave MacCallum, and Michael Weisberg for reading earlier drafts of this paper and providing me with some extremely helpful comments. In addition, I would like to thank the following people for some helpful correspondences on issues relating to neuroscience and physics: Peter Dayan, Stuart Hameroff, Christof Koch, Sebastian Seung, Henry Stapp, and Max Tegmark. Also, a version of this paper was read at a conference in Bled, Slovenia, and I would like to thank the members of the audience there for some very helpful feedback. Finally, the writing of this paper was partially funded by a year-long fellowship from the National Endowment for the Humanities and a quarter-long grant from California State University, Los Angeles; I would like to thank both of these institutions.

References

Balaguer, M. (2004). A coherent, naturalistic, and plausible formulation of libertarian free will. *Nous, 38*, 379–406. doi:10.1111/j.0029-4624.2004.00475.x.

Balaguer, M. (forthcoming). *Free will as an open scientific problem*. Cambridge, MA: MIT Press.

Bell, J. S. (1964). On the Einstein-Podolsy-Rosen paradox. *Physics, 1*, 195–200.

Bell, J. S. (1982). On the impossible pilot wave. *Foundations of Physics, 12*, 989–999. doi:10.1007/BF01889272.

Bohm, D. (1952). A suggested interpretation of quantum theory in terms of 'Hidden Variables'. *Physical Review, 85*, 166–193. doi:10.1103/PhysRev.85.166.

Bohm, D., & Hiley, B. J.(1993). *The undivided universe: An ontological interpretation of quantum mechanics*. London: Routledge.

Brasil-Neto, J. P. (1992). Focal transcranial magnetic stimulation and response bias in a forced choice task. *Journal of Neurology, Neurosurgery, and Psychiatry, 55*, 964–966. doi:10.1136/jnnp.55.10.964.

Dayan, P., & Abbott, L. F.(2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.

Deecke, L., Grotzinger, B., & Kornhuber, H. H. (1976). Voluntary finger movement in man: Cerebral potentials and theory. *Biological Cybernetics, 23*, 99. doi:10.1007/BF00336013.

Festinger, L. (1957). *A theory of cognitive dissonance*. Palo Alto: Stanford University Press.

Ginet, C. (1990). *On action*. Cambridge: Cambridge University Press.

Hagan, S., Hameroff, S. R., & Tuszyński, J. A. (2002). Quantum computation in brain microtubules: Decoherence and biological feasibility. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics, 65*, 061901. doi:10.1103/PhysRevE.65.061901.

Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research, 126*, 128–133. doi:10.1007/s002210050722.

Haggard, P., & Libet, B. (2001). Conscious intention and brain activity. *Journal of Consciousness Studies, 8*, 47–63.

Hameroff, S. R., & Penrose, R. (1996). Conscious events as orchestrated space-time selections. *Journal of Consciousness Studies, 3*, 36–53.

Hodgson, D. (2002). Quantum physics, consciousness, and free will, in Kane (2002), pp. 85–110.

Honderich, T. (1988). *A theory of determinism: The mind, neuroscience and life-hopes*. Oxford: Oxford University Press.

Honderich, T. (2002). Determinism as true, compatibilism and incompatibilism as false, and the real problem, in Kane (2002), pp. 461–476.

Isen, A., & Levin, P. (1972). Effect of feeling good on helping. *Journal of Personality and Social Psychology, 21*, 384–388. doi:10.1037/h0032317.

Kane, R. (1996). *The significance of free will*. New York: Oxford University Press.

Kane, R. ed. (2002). *The Oxford handbook of free will*. New York: Oxford University Press.

Kochen, S., & Specker, E. P. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics, 17*, 59–87.

Kornhuber, H., & Deecke, L. (1965). Hirnpotentialanderungen bei Willkurbewegungen und passiven Bewegungen des Menschen. *Pflugers Archiv fur die Gesamte Physiologie des Menschen und der Tiere, 284*, 1–17. doi:10.1007/BF00412364.

Libet, B. (2002). Do we have free will? in Kane (2002), 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. doi:10.1093/brain/106.3.623.

Mele, A. (2006). *Free will and luck*. 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. doi:10.1111/j.1475-4975.2005.00112.x.

Penrose, R., & Hameroff, S.R. (1995). What 'Gaps'? Reply to Grush and Churchland. *Journal of Consciousness Studies, 2*, 98–111.

Tegmark, M. (2000). The importance of quantum decoherence in brain processes. *Physical Review E: statistical physics, plasmas, fluids, and related interdisciplinary topics, 61*, 4194. doi:10.1103/PhysRevE.61.4194.

Velmans, M. (1991). Is human information processing conscious? *Behavioral and Brain Sciences, 14*, 651–669.

Walter, H. (2001). *Neurophilosophy of free will*. Cambridge, MA: MIT Press.

Wegner, D. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.