

A Coherent, Naturalistic, and Plausible Formulation of Libertarian Free Will

MARK BALAGUER

California State University, Los Angeles

1. Introduction

Let *libertarianism* be the view that human beings possess *L-freedom*, where *L-freedom* is defined as the ability to make decisions that are simultaneously (a) undetermined and (b) appropriately non-random. Much needs to be said about what appropriate non-randomness amounts to, but for now, let me just say that the central requirement that a decision needs to satisfy in order to count as appropriately non-random is that of having been *authored and controlled* by the person in question; i.e., it has to have been *her* decision, and she has to have controlled which option was chosen. Many (probably most) philosophers think that libertarianism is incoherent. The main reason for this is captured by the following little argument:

Any event that's undetermined is uncaused and, hence, accidental. That is, it *just happens*; i.e., happens randomly. Thus, if our decisions are undetermined, then they are random, and so they couldn't possibly be "appropriately non-random". Or to put the point the other way around, if our decisions are appropriately non-random, then they are authored and controlled by us; that is, *we determine* what we choose and what we don't choose, presumably for rational reasons. Thus, if our decisions are appropriately non-random, then they couldn't possibly be undetermined. Therefore, libertarianism is simply incoherent: it is not possible for a decision to be undetermined and appropriately non-random at the same time.

One might call this the Hobbes-Hume-Hobart argument.¹ It is possible to give a much more detailed and careful version of the argument, but this simple version will be good enough for our purposes; for I am going to

argue in this paper that any argument of this general kind (regardless of how the details go) is wrongheaded. I will do this by showing that appropriate non-randomness is not only consistent with, but virtually *entailed by*, the appropriate kind of indeterminacy. And I will do this without appealing to anything like Cartesian souls or irreducible agent causation; that is, I will assume that some form of mind-brain materialism is true and that all causation is ordinary event causation.

The libertarian view I will be defending against the above argument is similar in a number of ways to other libertarian views in the literature, but it is very original (one might even say radical) in one crucially important way. The similarities to other views will emerge as we proceed, but for now, let me say that the view I have in mind is *event-causal* (i.e., on this view, decisions are caused but not necessitated by prior events) and *non-Valerian* (i.e., on this view, the important indeterminacy is located at the moment of choice, rather than prior to choice, as with Valerian views). This sort of view was hinted at by Wiggins (1973) and later developed by Kane (1985, 1996, 1999), Nozick (1981), and Ekstrom (2000);² my view is most similar to Kane's because of the prominence he gives to what I will call *torn decisions* (which I will characterize in section 2.1). My view also bares some similarities to the *noncausal* non-Valerian views of Ginet (1990, 2002), McCann (1998), McCall (1994), and Goetz (1997), which take free decisions to be uncaused. In the end, the view I have in mind is more naturally described as event-causal, but as we'll see, it has "noncausalist leanings," and I think it goes some way toward blurring, or perhaps complicating, the event-causal/noncausal distinction. Finally, my view is less similar to Valerian views, which have been developed, though not endorsed, by Dennett (1978) and Mele (1995), and to agent-causal views, which have been developed by a number of philosophers, most recently by O'Connor (2000) and Clarke (1993).

But there is one important way in which my view differs from all of these other views. Almost everyone involved in this debate, whether libertarian or anti-libertarian, believes something along the following lines:

Look, it may very well be that our decisions are undetermined in some sense; that this is a real possibility arises from the fact that the standard (and probably best) interpretations of quantum mechanics take the world to be indeterministic at the most fundamental level. It may be that the world is causally indeterministic all over the place, and so it may be that neural events and, hence, decisions are indeterministic. So *half* of the libertarian thesis is a real possibility. But the other half—appropriate non-randomness—is the really hard part.

Again, just about everyone thinks that something like this is right. The difference between libertarians and anti-libertarians is that the former think that while it is indeed very hard to see how indeterminism and

appropriate non-randomness could be compatible, in the end, they *are* compatible, and so libertarianism is coherent; whereas the latter think that if our decisions are undetermined, then they couldn't be appropriately non-random. But I reject all of this. In section 3, I will argue for the following claim:

(* There is a significant subset of our decisions (in particular, *torn* decisions, which again, will be defined in section 2.1) for which the following is true: if they are undetermined (at the moment of choice), then they are appropriately non-random.

Notice that I am not just saying that if our torn decisions are undetermined, then they still *could* be appropriately non-random, so that libertarianism is coherent. There would be nothing new in that; Kane (1985, 1996) and Ginet (1990) have argued precisely this point, and indeed, just about every libertarian (certainly all those mentioned in the preceding paragraph) has tried at least to defend the thesis that libertarianism is coherent. In contrast to this, I am going to offer a positive argument for the much stronger thesis that if our torn decisions are undetermined at the moment of choice, then they *are* appropriately non-random. Thus, on my view, it is indeterminacy and not appropriate non-randomness that is the really controversial thesis in libertarianism. Indeed, if I'm right about (*), then the question of whether libertarianism is true just reduces to the question of whether some of our torn decisions are undetermined at the moment of choice.³ This, of course, is a straightforward empirical question about the neural events that are our torn decisions. Moreover, it is a *wide open* question; in particular, as of right now, we have no reason whatsoever to believe that all of our torn decisions are determined, and hence, no reason to doubt that libertarianism is true. I don't have the space to argue this point here, but I will say a few words about it in section 5. Finally, since a well-developed libertarianism should provide some account of decisions that aren't torn, I will say a bit in section 4 about how various kinds of non-torn decisions fit into the overall libertarian view that I have in mind.

Before going on, I should say a few words about the kind of mind-brain materialism that I will assume. In short, I want to assume only a very weak version of materialism. Once we reject mind-brain dualism, we are left with the view that decisions are physical events, presumably neural events; the weakest version of this view is the *token-token identity theory of decisions*, and this is all I will assume here. The token-token theory says that every particular (human) decision is a particular neural event, but it does not take a stand on the type-type question, i.e., the question of whether all decisions are neural events of some specific kind. It is worth emphasizing how weak and uncontroversial the token-token theory is: it is consistent with—indeed, entailed by—all of the standard (non-eliminativist) views in the philosophy of mind, e.g., type-type identity theories, functionalist theories, non-reductive

physicalist theories, and so on; but on the other hand, the token-token theory doesn't entail any of these other views; all it says is that decisions are physical events, presumably neural events. I think it's safe to say that we have very good reasons to believe this theory and that, at present, almost everyone working in psychology, neurophysiology, and the philosophy of mind accepts it. (The only other tenable view that an anti-dualist could hold here is the eliminativist view that there are no such things as decisions; I will ignore this view here, not just because it is implausible or because so few people endorse it, but because my argument can be run in eliminativist terms as well as token-token-identity terms, so that in the present context, it just doesn't matter whether eliminativism about decisions is true.⁴)

Finally, I should point out that I will not be concerned here with the question of whether the notion of L-freedom captures the ordinary-language notion of *free will*, or "real free will", or the free will we care about. These issues arise in connection with the compatibilism debate, but I won't discuss them here because they are not relevant to the questions that I'm concerned with in this paper, namely, the question of whether L-freedom is a coherent notion and the question of whether human beings actually possess L-freedom, i.e., whether libertarianism is true. (Some philosophers use 'libertarianism' to denote the view that human beings possess L-freedom *and* that this is the proper way to understand the ordinary notion of free will. That's not how I'm using the term here, although of course, nothing important hangs on this terminological point.)

2. Preliminaries

2.1. Torn Decisions

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—i.e., the person has the experience of "just choosing". (In such cases, subconscious reasons, and subconscious weightings of reasons, can come into play; we will see later that these do not alter my argument, but for now, I want to ignore the complication of the subconscious.)

Let's look at two examples of torn decisions. First, let us suppose that Ralph, a life-long resident of Mayberry, North Carolina, is trying to decide whether or not to move to New York City. He has safety-and-stability-based reasons for wanting to stay in Mayberry (he's been offered a position as assistant day-shift manager at the local Der Wienerschnitzel and is clearly being groomed for the manager position, and his sweetheart, Robbi Anna, has offered him her hand in marriage); and he has fame-and-fortune-based reasons for wanting to move to New York (he longs to be the first person to start at middle linebacker for the Giants while simultaneously starring on

Broadway in a musical production of Sartre's *Nausea*). He deliberates for several days, considering all of his reasons for choosing, but he is unable to come to a view as to which set of reasons is stronger. He feels genuinely torn. But Robbi Anna and Der Wienerschnitzel, Inc., are waiting, and he feels he has to make a decision. (He may actually have a deadline, but this isn't necessary for the example.) Finally, he decides to move to New York, but not because he comes to believe that his reasons for moving to New York outweigh his reasons for staying in Mayberry. He was unable to come to a view either way on that question, and in the end, he *just decided to go*. Period.

Torn decisions aren't always "big, life decisions," and indeed, it seems to me that we make decisions like this all the time, every day of our lives. To appreciate this, consider a second case, involving a decision more mundane than Ralph's. Jane is in a restaurant, deliberating about whether to order tiramisu or a fruit plate for dessert. She thinks that the former will taste better but that the latter will be better for her health. She has no clue which reason is stronger and feels genuinely torn. Suddenly, it's her turn to order; the waiter is looking at her; she has to pick; Oh, God; "I'll have the tiramisu," she says. I.e., she just chose. Period.

Of course, this is just a description of how torn decisions feel to us: it *seems* that we just make them. For all we know, however, it may be that all of our torn decisions are determined, e.g., by physical events in our brains that precede these decisions. Whether or not our torn decisions are in fact causally determined is an empirical question.⁵

2.2. Indeterminacy

As I pointed out above, when I say that if a torn decision is undetermined, then it is appropriately non-random, what I mean is that if it's undetermined *at the moment of choice*, then it's appropriately non-random. Thus, if I say that a decision was determined, that does not mean it was determined prior to the agent's birth, or any such thing; if an undetermined quantum event occurs in my head two seconds prior to a decision, and if this event (together with physical laws and other circumstances) causally determines my decision, then on my usage, the decision counts as determined. Prior-to-choice indeterminacies might be relevant to the decision in a variety of ways, but as we'll see, they aren't needed for L-freedom. This is what makes my libertarian view non-Valerian.

A second point that needs to be made here is that an agent's reasons could causally determine that a given torn decision will occur, and could causally determine that the choice will come from a specific set of possible options, without determining which option will be chosen. E.g., it could be that (a) Jane had compelling reasons for choosing *something* when it was her turn to order, and these reasons *caused* her to just pick; and (b) it was causally determined by Jane's reasons that she would choose either the

tiramisu or the fruit plate, i.e., that she wouldn't pass on desert or order something else, e.g., pea soup; but (c) the fact that Jane chose the tiramisu over the fruit plate was not causally determined (by her reasons or anything else). What I mean when I say that a torn decision was undetermined at the moment of choice is captured by clause (c): it wasn't determined which option was chosen. Whether or not it was causally determined that *some* choice would be made, or that the choice would come from some specific set of options, will not matter.

A third issue that needs to be discussed here is this: whatever actually happens in human heads leading up to torn decisions, conceptually speaking, there can be different *degrees of determinacy* here. Let me explain what I mean by this. In any torn decision, the agent's reasons are neutral between a set of tied-for-best options. Thus, the reasons-based probabilities of the various live options are equal (or at least roughly equal). But it could be that factors external to the agent's reasons (e.g., wholly non-mental neural events) causally influence the choice and wholly or partially determine which option is chosen. Indeed, given that we are allowing for probabilistic (i.e., non-necessitating) causation, there is a continuum of possible cases here. At one end of the spectrum, which option is chosen is wholly uncaused (though, again, it may be that it was causally determined that the agent would choose *something* and, indeed, that she would choose from some specific set of options); at the other end of the spectrum, which option is chosen is causally determined, i.e., nomologically necessitated by the state of the universe together with causal laws; and in between, there is a continuum of possible cases, where factors external to the agent's reasons make the moment-of-choice probability of some specific option being chosen higher and higher. I think that libertarians should respond to this by allowing that (conceptually speaking) there can be different *degrees of L-freedom*. Combining this with the thesis of the present paper (i.e., the thesis that if our torn decisions are undetermined, then they're appropriately non-random and L-free), we are led to the following two-part thesis:

(A) If a torn decision is "wholly undetermined" at the moment of choice (i.e., if which option is chosen is uncaused, so that the moment-of-choice probabilities of the various reasons-based tied-for-best options are at least roughly even, given the complete state of the world and all the laws of nature), then the decision is wholly appropriately non-random and wholly L-free. And

(B) Our torn decisions are appropriately non-random and L-free to the extent that they are undetermined at the moment of choice; thus, to take the example of Ralph, there is a continuum of conceptually possible cases, where factors external to Ralph's reasons make his choosing to move to New York more and more likely and, hence, the decision becomes more and more determined; the claim here is that as the decision becomes more

determined by factors external to Ralph's reasons, it becomes less appropriately non-random and, hence, less L-free.

I think this is the right way for libertarians to conceptualize things, but I will not have the space here to argue for both (A) and (B), and so I am just going to argue for (A). Now, I think that (A) is clearly the harder of the two theses to motivate, and indeed, I think it will be clear from my argument how I would go on to argue for (B). But in this paper, I will argue for (A) only, and so the relevant kind of indeterminacy here is the one at work in (A)—i.e., the one captured by the idea of a *wholly* undetermined event. Thus, when I say that if our torn decisions are undetermined at the moment of choice then they are appropriately non-random, what I mean is this: if our torn decisions are wholly undetermined in the sense of (A)—i.e., if they are such that, at the moment of choice, which option is chosen is uncaused, so that the moment-of-choice probabilities of the various reasons-based tied-for-best options are at least roughly even, given the complete state of the world and all the laws of nature—then these decisions are appropriately non-random. For the sake of brevity and clarity, though, I will often use 'undetermined at the moment of choice', as opposed to this mouthful of words. It is a bit non-standard to use 'undetermined' in this way, for usually, it is used to mean 'not strictly determined', i.e., 'not causally necessitated'; but given that there can be a continuum of different degrees of causal determination, and given that I am arguing here for thesis (A), which is concerned only with the limiting case of wholly undetermined decisions, I think that this is the clearest way for me to speak.

I should also note here that by proceeding in this way, I do not make my libertarian view implausible, for as I will point out in section 5, as of right now, there is no good reason to doubt the thesis that for at least some of our torn decisions, which option is chosen is uncaused at the moment of choice. That is, there is no good reason to believe the thesis that all of our torn decisions are, at the moment of choice, significantly causally influenced by factors external to us. Indeed, there is no more reason to believe this thesis than there is to believe the thesis that all of our torn decisions are strictly determined at the moment of choice, and so by proceeding in this way, I have not made my libertarian view implausible.

(It's also worth noting that on the (A)–(B) view, the question of whether libertarianism is true is a question about the degree to which human beings are L-free, and so this question reduces not just to the question of whether some of our torn decisions are wholly undetermined at the moment of choice, but also to the more complicated (but still purely empirical and wide-open) question of the extent to which our torn decisions are causally determined at the moment of choice.)

2.3. *Appropriate Non-Randomness*

As I said above—and I am here just echoing the standard view of the matter—the heart of appropriate non-randomness is captured by the idea that in order for a decision to be L-free, it has to be authored and controlled by the person in question; i.e., it has to be *her* decision, and she has to control which option is chosen. But while authorship and control form the heart of appropriate non-randomness, this is not all there is to the story. At least in connection with torn decisions, appropriate non-randomness also requires what Kane has called *plural authorship and control*,⁶ which can be understood as follows: if an agent S is trying to decide between multiple alternatives, say A, B, and C, and eventually chooses A, then in order for the choice to be L-free, it must be the case not just that (i) the choice is undetermined, i.e., that S could have chosen some other option (e.g., B) instead of A, and that (ii) the choice is authored and controlled by S, but also that (iii) there is at least one other option, say B, such that if S had chosen B, then this choice would also have been authored and controlled by S. Moreover, appropriate non-randomness also seems to require some sort of *rationality* (and at least in connection with torn decisions, *plural rationality*); in section 3.2, I will describe a few different kinds of rationality, and I will try to determine which of them are required for appropriate non-randomness and L-freedom.

3. The Argument

In this section, I will argue that

(**) For any ordinary human torn decision, if it is undetermined (i.e., if it is such that, at the moment of choice, which option is chosen is uncaused, so that the moment-of-choice probabilities of the various reasons-based tied-for-best options are at least roughly even, given the complete state of the world and all the laws of nature), then it is also appropriately non-random (i.e., it is authored and controlled by the agent in question in an appropriately rational way and in a way that satisfies the requirement of plural authorship, control, and rationality) and hence, it is L-free.

I will argue for this thesis by addressing the three requirements for appropriate non-randomness in turn; I will discuss authorship and control in section 3.1, rationality in section 3.2, and plural authorship, control, and rationality in section 3.3.

3.1. *Authorship and Control*

In this subsection, I will argue that if our ordinary torn decisions are undetermined at the moment of choice, then we author and control

them. One of the keys to my argument here is the token-token identity theory of decisions. This is a bit ironic, because mind-brain materialism has often led people to doubt libertarianism. People sometimes say things like this:

If decisions are just neural events, and if in addition, they're not determined by the neural states that are our reasons (i.e., our beliefs, desires, and so on), then presumably they arise out of brutally physical (i.e., wholly non-mental) neural events in our brains, and so it could not be that we author and control them.

But this is confused. This is what we would say to a materialist who claimed that our decisions are determined but not by our reasons. But that's not the picture I'm proposing at all. On the view I have in mind, our torn decisions aren't determined by *anything*: they're not determined by our reasons, and they're not determined by the physical states of our brains just prior to the moment of choice; rather, we just make them. If the token-token identity theory of decisions is right, and if our torn decisions are undetermined at the moment of choice in the way described above, then we ought not to say that these decisions "arise out of" neural events; what we ought to say is that they *are* neural events (and that they are undetermined). But while they are neural events, if the token-token identity theory is true, they are also *mental* events; in particular, they are purposeful, intentional decisions; Ralph, for instance, decided to move to New York consciously and on purpose, fully aware of what he was doing and in full possession of his faculties. Thus, it doesn't follow from the mere fact that our torn decisions are physical events that they're unfree, or random in some appropriate sense; for given the token-token identity theory, our torn decisions are still *mental* events, and in particular, they are still conscious, intentional, purposeful, and so on.

This opens the door to an initial argument for the thesis that if our torn decisions are undetermined at the moment of choice, then we author and control them. The first point to note here is that if a torn decision is undetermined at the moment of choice in the way described above, then it follows that nothing external to the agent causes her to choose as she does. (E.g., if Ralph's decision was undetermined at the moment of choice, then nothing external to him made him choose as he did. Quantum particles and the laws of physics didn't make him do this; an alien intervener didn't make him do it; *nothing* made him do it.) And the second point is that when we combine this lack of external causation with conscious, intentional, purposefulness, we seem to get authorship and control: if (a) an agent S consciously, intentionally, and purposefully chooses some option A, and (b) nothing external to S causes her to choose A, then it seems that (c) she authors and controls the decision. E.g., if we assume that Ralph's decision was undetermined at the moment of choice, then we have that (a) Ralph

chose consciously, intentionally, purposefully, and so on, and (b) nothing external to him caused him to choose as he did. So why *wouldn't* we say that he authored and controlled his decision? What more could we want out of authorship and control than this—i.e., than (a) and (b)?

Well, one might think that more *is* needed for authorship and control. Indeed, I can think of two different ways in which one might try to argue this point, giving rise to two objections to this initial argument for the thesis that if our torn decisions are undetermined at the moment of choice, then we author and control them. The objections are as follows:

Objection 1: If a decision is undetermined, then it isn't determined by the agent's reasons, or character; but then at the very least, the agent has less control and authorship than she would have had if her decision had been fully determined by her reasons and character. In short, the looser the connection between the agent's reasons and character on the one hand and her decision on the other, the less authorship and control she has over the decision.

Objection 2 (a.k.a. the "luck objection", or the "chance objection"): If a decision is undetermined at the moment of choice, then given the same past, the agent could just as easily have made another choice without anything about the agent changing, and so the agent could not be the source of, or have control over, the choice. To bring this out, suppose that some agent S is torn between two options, A and B, and eventually chooses A in a torn-decision sort of way. And now suppose that God "replayed" the universe 99 times (or "rolled back" the universe 99 times) and that in 50 of the "plays" of the universe, S chose A, whereas in the other 50 plays, she chose B. Finally, suppose that God replays the universe again, and that we are waiting to see what S will choose this time. Isn't it just a matter of *chance*, or *luck*, what S will choose? And so isn't it also a matter of chance or luck what S chooses on the *first* play of the universe? Another way to appreciate this point is to notice that even if prior conditions don't determine whether S will choose A or B, it might still be that prior conditions determine that there are certain fixed probabilities of A and B being chosen (e.g., in the case of S, it seems that both probabilities are around .5); but given this, isn't it just a matter of chance or luck whether, in fact, S chooses A or B? And of course, if it's a matter of chance or luck, then it seems that S doesn't author or control the decision. (Arguments of this general kind have been mounted in a number of different ways in recent years by a number of different philosophers, including van Inwagen (2002), Fischer (1999), Mele (1999a,b), Haji (1999), Clarke (1995, 2002), O'Connor (2000), Double (1991), Strawson (2000), Waller (1988), Bernstein (1995), and Berofsky (2000),⁷ many of these arguments, I should say, have been directed against Kane's version of libertarianism.)

Let me respond first to objection 2. If it's really true that if we "replayed" the universe several times, Ralph would opt for New York about half the time and Mayberry about half the time, then how can we say that Ralph authored and controlled his decision to move to New York? Isn't it just a matter of chance, or luck, that he made that choice in the one real "play" of the universe?

No, it's not *just* a matter of luck or chance. It's true that, in some sense, Ralph's decision was arbitrary (or if you like, chancy or lucky). But this chancy event was still a *decision* (it was conscious, intentional, purposeful, and so on) and it was still *Ralph's* decision: nothing external to him made him choose as he did. So it's not obvious that Ralph didn't author and control the decision. This is a delicate matter. The question is whether the intuitive notion of *being in control of what you do* applies in this case. I want to argue that if Ralph's decision was undetermined at the moment of choice, then the intuitive notion of control clearly does apply (and that the notion of authorship does as well), despite the fact that the choice was, in some sense, arbitrary or chancy or lucky.

I begin my argument by pointing out that because Ralph's decision was a *torn* decision, there is no view that procures more authorship and control for his decision than my libertarian view procures. The fact of the matter is that Ralph's reasons do not pick out a best option; but he does have reason to make *some* choice, i.e., for not remaining in a state of indecision; thus, it seems clear that Ralph has good reason to *just pick* from among his tied-for-best options (i.e., to pick randomly, or arbitrarily, between staying in Mayberry and moving to New York). Given this, it seems that the most we could hope for here, *vis-à-vis* authorship and control, is that it be *Ralph* who does the just-picking, i.e., who makes the random selection from among his tied-for-best options. If the selection is made by an alien intruder, or by a non-mental brain event in Ralph's head that triggers the choice, then it seems clear that in some important sense, Ralph does not author or control the decision. But if the mental/physical event that is Ralph's decision is undetermined at the moment of choice in the way we're assuming here, then nothing external to Ralph causally influences his choice, and so it seems that in this case, we do get the result that it is Ralph who does the just-choosing. For in this case, the *only* event that is relevant at all to whether Ralph chooses to move to New York is the decision itself; but the decision itself is a *mental* event—it is conscious, intentional, purposeful, and so on—and it is *Ralph's* mental event. And so it seems that in this scenario we procure as much authorship and control for Ralph as we can, given that he is making a torn decision.

To further appreciate this point, suppose that God told us that we could dictate how the world would be with respect to our torn decisions, and suppose that in deciding how to set things up, the only thing we cared about was maximizing authorship and control. How ought we to proceed? Well, it

seems to me that we couldn't ask for anything better than a scenario in which the agent gets to just-choose from among her tied-for-best options, without anything causing her to pick one way or the other. For (a) if the choice were anything other than a *random* selection—i.e., anything other than a just-choosing—then the selection procedure wouldn't mesh with the agent's reasons, because her reasons are, by assumption, neutral between her tied-for-best options; and (b) if the just-choosing were done by anything other than the *agent*, then she would lose control.

In response to this, one might say something like the following: "It would be better to set things up so that people would never be forced to *just choose*, i.e., so that people would always have good reasons for their choices." I have three things to say about this. First, to set things up like this is not to set them up so that we have more control over our torn decisions; rather, it's to set them up so that we never make torn decisions. Second, it's hard to see why anyone would *want* to set things up like this, because there doesn't seem to be any reason for thinking it undesirable to have to make torn decisions, i.e., for thinking that lives without such choices are better, *ceteris paribus*, than lives with them. Indeed, my intuition goes the opposite way: if we never made any torn decisions, if we always had compelling reasons for all of our choices, then we could be spontaneous only by being irrational, and that, I think, would be bad; we would be slaves to our reasons, and so intuitively, we would have *less* control than we would if we made some torn decisions in which our choices weren't causally influenced by anything external to us. Third and most important, by setting things up so that we never made any torn decisions, we wouldn't be increasing authorship or control at all, because in fact, we do not have more authorship or control over choices that are uniquely picked out by our reasons than we do over torn decisions that are undetermined at the moment of choice. To appreciate this, consider two Ralphys. $Ralph_1$ is the Ralph we've been considering so far, i.e., the Ralph who is torn between his two options and in the end arbitrarily chooses (i.e., has the experience of *just choosing*) to move to New York, without anything external to him causally influencing his choice. $Ralph_2$, on the other hand, comes to believe that his reasons for moving to New York are stronger than his reasons for staying in Mayberry and, therefore, decides to move to New York. Does $Ralph_2$ have more authorship or control over his decision than $Ralph_1$? I find it hard to see why we should say that he does. For consider: $Ralph_1$'s decision was a mental event (it was conscious, intentional, purposeful, etc.), and nothing external to him made him choose as he did. So it is *Ralph₁*, and nothing else, who just chose to move to New York. Again, it's hard to see why $Ralph_1$ has any less authorship or control than $Ralph_2$. The difference between the two Ralphys isn't that $Ralph_2$ has more control— $Ralph_1$, after all, can do whatever he wants to do; it's simply that $Ralph_2$ has compelling reasons for his choice, whereas $Ralph_1$ does not.

It is worth noting that I am disagreeing here with something Kane (1999) has said on this topic. He says that indeterminism diminishes a certain kind of control while increasing another kind. In contrast, I think that (i) given that a decision is a *torn* decision, the sort of indeterminism I've been describing is the only way to guarantee control for the agent; for since the agent's reasons don't pick out a unique best option, anything that does determine this (e.g., a non-mental brain event in the agent's head) would presumably be extra-agential, and so if our torn decisions are determined, then they are not under our control (or at the very least, we can say that they *might not* be under our control, because they might be determined by events of the wrong kind, and I don't need anything stronger than this here⁸); and (ii) for the reasons given in the preceding paragraph, if our torn decisions are undetermined at the moment of choice, then we have *just as much* control over them as we do over non-torn decisions in which our reasons determine our choices.⁹

It is also worth noting here that point (i) from the preceding paragraph gives me an argument against something Clarke (1995, 2002) has said, namely, that indeterminism doesn't increase control. If point (i) is correct, then for torn decisions, indeterminism does increase control in the sense that it *procures* control; for if our torn decisions were determined, then we wouldn't (or at the very least, might not) control them. (Clarke thinks that in order to procure control, indeterminism is not enough; we also need agent causation. But I don't see how this is supposed to help. For (a) my libertarian view already delivers the result that, in torn decisions, the agent chooses without being causally influenced by anything external to her (and as we'll see below, it also delivers the result that in both torn and non-torn decision, the agent chooses for her reasons); and (b) I don't see how we could do any better than this by introducing a notion of agent causation. Now, I should say that my libertarian view does allow for a certain kind of agent causation, but this will be reducible to event causation—in particular, to causation by the agent's decisions—and so according to standard usage, my libertarian view is not an agent-causal view.¹⁰)

One last point about objection 2: some people, e.g., Fischer (1999), seem to think that the mere fact that Ralph would choose differently in multiple "plays" of the universe already suggests that he is not in control. I find this hard to understand. If in each play of the universe, it is *Ralph* who does the just-picking, then doesn't he control the choice? If Ralph chooses to stay in Mayberry in play number 2, why would that undermine the fact that it was *him* who chose to move to New York in play number 1? Indeed, I think just the *opposite* is true; given that Ralph's reasons don't pick out a unique best option, I would find it very suspicious if he always chose the same way in multiple plays of the universe. If he always chose to move to New York, then I would think, "Look, that can't be a coincidence; something must be *causing* him to choose in that way; and since (by assumption) his reasons aren't causing this, it must be something else, e.g., a random, non-mental

event in his nervous system.” So I would say that given that Ralph is making a *torn* decision, what we would *expect*, if the decision is flowing from *him*, is that he would choose differently in different plays of the universe. And so again, this supports the idea that, in torn decisions, indeterminism of the sort I have in mind procures control.

Let me move on, finally, to objection 1. The central premise of this objection is the following:

Principle D: Determination by reasons is required for full authorship and control, and in general, the looser the connection between the agent’s reasons and character on the one hand and her decision on the other, the less authorship and control she has over the decision.

Now, in the present context, it would clearly be question-begging to take this principle as a basic premise without argument. So the question is whether there is any argument in its favor, and I must confess that I can’t imagine what such an argument might look like. I admit that if we focus only on cases in which the agent’s reasons pick out a unique best option, then Principle D can seem intuitively plausible (and indeed, my view allows for the claim that in such cases Principle D is true, and, as we’ll see in section 4, it allows for the claim that such decisions are L-free). But as soon as we think about torn decisions, Principle D seems completely implausible. For it entails that regardless of whether torn decisions are determined or undetermined, and regardless of the overall view of free will that we end up endorsing, no human being has ever authored or controlled a torn decision. That strikes me as wildly implausible. If Sandy has to choose between A and B, and if her reasons don’t pick out a unique best option, it simply doesn’t follow that for this reason alone, she can’t make a *decision* at all, or that when a decision does happen, it isn’t *Sandy* who makes and controls the decision. Intuitively, the notions of authorship and control seem to allow for a random selection here, so long as it is *Sandy* who makes the random selection and not something external to her.

In fact, I have already argued that Principle D is false; for I have argued that an agent can have as much control over a torn decision that’s undetermined at the moment of choice as she does over a choice that is determined by her reasons. So in the absence of an argument to the contrary, I am inclined to conclude that in connection with torn decisions, Principle D just fails. (I should note, however, that even though Principle D fails for torn decisions, there might still be a weakened kind of determination by reasons here. In particular, it could be that when we make torn decisions, our reasons cause us (i.e., causally determine us) to choose from among our tied-for-best options. And indeed, this seems plausible; e.g., it is plausible to suppose that it was causally determined that Ralph was not going to decide to move to Outer Mongolia—i.e., that he was going to choose between New York and Mayberry. And this, I think, dovetails with the claim that the decision was *Ralph’s*.)

My conclusion, then, is that if our torn decisions are undetermined at the moment of choice, then we author and control them. And I would like to point out here that this conclusion is introspectively satisfying. To appreciate this, suppose that a race of super-intelligent neuro-cognitive scientists studied your brain and told you that, in fact, your own torn decisions are undetermined at the moment of choice, so that when you make these decisions, nothing external to you makes you choose as you do. Would you conclude from this that you do not author or control these decisions? It seems to me that it would be downright bizarre to draw that conclusion. Indeed, if I found out that when I make my torn decisions, nothing external to me causally influences how I choose, I would conclude from this (together with what I already know about these decisions, namely, that they’re conscious, purposeful, intentional, and so on) that I *do* author and control these decisions. And this suggests to me that the intuitive notions of authorship and control do apply here, and hence, that if my torn decisions are undetermined at the moment of choice, then I author and control them.

Before going on, I want to do two things. First, I want to say something about a theory of authorship and control that seems to be inherent in what I have said here. Consider the following:

(A&C) An agent A has authorship and control with respect to a decision D if (i) A makes D in a conscious, intentional, purposeful way, and (ii) nothing external to A and her reasons causes her to choose as she does. (And if something external to A is probabilistically causally relevant to how A chooses—i.e., if it increases the probability of some option being chosen, without necessitating it—then A’s authorship and control are diminished to the extent that the given probability is increased by the extra-agential factor.)

One might object to this theory on the grounds that if an alien implanted reasons into my head, and I acted on them, I wouldn’t be in control. I might quibble over whether (ii) is really satisfied here, but I have a different response: I don’t intend (A&C) as a conceptual analysis (i.e., as providing a sufficient condition for the concepts of authorship and control); rather, I intend it as a claim about when *actual* humans have authorship and control. Thus, since it’s reasonable to suppose that in fact, no aliens have implanted any reasons in our heads, this objection is irrelevant. And indeed, this is how I intend the entire paper to be read: I am claiming that if *our* torn decisions are, *in fact*, undetermined at the moment of choice in the way described above, then they are, in fact, appropriately non-random.

Finally, I want to say a few words about subconscious reasons. Let’s return to Jane’s decision and suppose that (a) consciously, Jane believes that her reasons for choosing do not favor either of her options, but (b) subconsciously,

she hates her mother and believes that ordering tiramisu would be a good way of getting back at her mother because her mother always detested tiramisu, and (c) these subconscious states cause her to choose tiramisu, although she experiences this as “just choosing”. The point I want to make about this sort of case, in the context of the present argument, is that it is simply *not a case of a decision that's undetermined at the moment of choice*. Jane's decision here is causally determined by subconscious mental states. Therefore, this sort of case is simply irrelevant to the argument that I am mounting in this section, because I am trying to argue only that *if* our torn decisions are undetermined at the moment of choice, then they're appropriately non-random. (There could also be cases in which subconscious reasons increase the probability of some option being chosen, without causally determining it. In such cases, libertarians can say that the agent's authorship and control are diminished to the extent that the subconscious reasons increase the probability of the given option being chosen. Or they might want to argue that (deterministic or probabilistic) causation by subconscious reasons actually *doesn't* diminish control, or perhaps diminishes it less than causation by wholly extra-agential factors. I don't need to take a stand on this issue, because it is independent of the thesis of this paper.)

We can say the same thing about decisions that are caused by subconscious weightings of reasons, or by compulsions. If Jane subconsciously gave more weight to her tiramisu reasons (or if she is a compulsive tiramisu eater) and if this caused her to choose as she did, then her decision was not undetermined at the moment of choice, and so it is not relevant to my argument.

One might object here as follows: “You say that cases of this sort are irrelevant to your argument, because you're trying to argue only that *if* our torn decisions are undetermined, then they're appropriately non-random. But it's an empirical question whether there even *are* any decisions that are genuinely undetermined at the moment of choice. Maybe there are no such things.” My response: of course there may be no such things; this is an empirical question. Now, in section 5, I will suggest that it is also a *wide-open* question and, in particular, that as of right now, we have no reason whatsoever to think that all of our torn decisions are determined. But it is also worth noting here that the issue of subconscious reasons (or subconscious weightings of reasons, or compulsions) doesn't itself provide a reason to think that all of our torn decisions are determined; of course, it *may* be that all our torn decisions are in fact determined by such things, but the important point here is that this is a wildly speculative empirical hypothesis and that, as of right now, we have no reason to believe it. (Notice how strong this hypothesis is: it entails that whenever we feel torn, we *always* either have subconscious mental states waiting there to settle the matter or else such states are generated on the spot. I can't imagine why anyone would believe this at the present time.)

3.2. Rationality

I have spent so much time on authorship and control because I think they are what's most important here. But I imagine that some readers will grant that on the model I'm constructing, Ralph and Jane do author and control their decisions but maintain that I have secured this result only at the cost of making it impossible to maintain that their decisions are rational. For instance, Dennett (1978) argues that Valerian libertarianism is superior to non-Valerian libertarianism by claiming that if the undetermined event occurs after deliberation, then it makes deliberation otiose. And Double (1991) argues that even if non-Valerian libertarian views can salvage control, or authorship, they cannot salvage rationality. In the present subsection, I will dispel these worries by arguing that if our torn decisions are undetermined at the moment of choice, then they are appropriately rational, i.e., rational in all the ways that are required for appropriate non-randomness and, hence, for L-freedom.

The first thing we need to do is get clear on what the appropriate senses of rationality are. To this end, let me define a few different ways in which a decision can be rational. First, I will say that a decision is *strongly rational* (given the agent's reasons for choosing) if and only if the agent's (conscious) reasons for choosing entail that the option chosen is the best available option. Second, a decision is *weakly rational* (given the agent's reasons for choosing) if and only if it is consistent with the agent's (conscious) reasons for choosing.¹¹ (Notice that weak rationality is certainly a kind of rationality; in particular, a decision that's weakly rational is not *irrational* in the sense of going against the agent's (conscious) reasons for choosing.) Third, a decision is *teleologically rational* if and only if it is made for (conscious) reasons. Thus, while weak rationality is opposed to a certain kind of irrationality, teleological rationality is opposed to a kind of arationality. (We could get more fine-grained in our distinctions here by introducing subconscious reasons. For instance, there could be a decision that was consistent with the given agent's conscious reasons but not her subconscious reasons; or a decision that wasn't made for any conscious reason but was generated by a subconscious reason. We will see below that the introduction of subconscious reasons does not alter my argument, but for now, I want to ignore this issue.)

Torn decisions are, by definition, weakly rational but not strongly rational: the agent chooses a reasons-based tied-for-best option, and so the decision is consistent with, but not entailed by, the agent's reasons. Torn decisions are also teleologically rational in at least two different senses (and I should note here that Kane (1985, 1996) has also argued that torn decisions are made for reasons). First, our reasons (and deliberations) serve at least to limit the set of options from which we end up choosing in our torn decisions. E.g., Ralph didn't decide to move to Outer Mongolia, and there's a very good explanation for this: he had no reason for so deciding; he

had no desire to move to Outer Mongolia. Moreover, it's plausible to suppose that reasons play a causal role here. For (a) it is just a fact that our torn decisions do dovetail with our reasons—in particular, when we make torn decisions, we almost always end up choosing one of our reasons-based tied-for-best options—and (b) it's hard to believe that this is a coincidence, that our reasons don't play a role in making this the case. Perhaps our reasons play a role in causally determining a set of possible options from which we choose; but however it happens, it seems hard to deny that our reasons are causally relevant to the genesis of our torn decisions, even if they don't causally determine precisely which option is chosen. And so it seems safe to conclude that our torn decisions are teleologically rational in at least this sense. Second, ordinary torn decisions like Ralph's are also teleologically rational in the sense that they are made *in order to satisfy reasons*. Ralph, for instance, decided to move to New York in order to try to become an actor and a football player. He might not have thought that these reasons were compelling reasons for moving to New York, but it's clear that he decided to move to New York in order to try to satisfy these desires of his.

So torn decisions are weakly and teleologically rational but not strongly rational. Thus, one point I need to argue here is that strong rationality is not required for appropriate non-randomness or L-freedom. We can appreciate this point very easily. Suppose someone said the following:

Jane did not have compelling reasons that uniquely picked out her choice of tiramisu, so she did not choose the tiramisu of her own free will, because her choice was not sufficiently rational.

I think it's pretty clear that this remark would seem downright bizarre to just about everyone. In our ordinary discourse and thought, we simply don't think that because decisions like Jane's aren't strongly rational, they are not free. Thus, strong rationality is not required for the ordinary notion of free will. Now, I suppose one might think that, despite this, libertarians are committed to the claim that strong rationality is required for L-freedom. But this is just false. We need to keep in mind that libertarians get to dictate what L-freedom requires; after all, it's *their* view. But no libertarian would sit still for the above indented comment about Jane's decision, so I think it's pretty clear that strong rationality is not required for L-freedom (or for the ordinary notion of free will).

At this point, we can bring in subconscious reasons and see that they present no problem for my thesis either. Once we acknowledge subconscious reasons, we can define more kinds of rationality. For instance, corresponding to weak, strong, and teleological rationality, we can define three more kinds of rationality as follows: a decision is *overall-teleologically rational* if and only if it is made for a (conscious or subconscious) reason; a decision is

overall-weakly rational (given the agent's reasons for choosing) if and only if it is consistent with the agent's (conscious and subconscious) reasons for choosing; and a decision is *overall-strongly rational* (given the agent's reasons for choosing) if and only if the agent's (conscious and subconscious) reasons for choosing entail that the option chosen is the best available option.¹² But (a) teleological rationality entails overall-teleological rationality, and so if our torn decisions have the former sort of rationality, they also have the latter; and (b) overall-strong rationality is clearly not required for L-freedom, or the ordinary notion of free will: in order for Ralph's decision to count as free, or L-free, it's clearly not required that he had subconscious reasons that entailed that moving to New York was his best option; and likewise, (c) overall-weak rationality isn't required either: even if Ralph had subconscious reasons for staying in Mayberry, so that overall, his conscious and subconscious reasons favored Mayberry over New York, it seems clear that since he believed that his reasons were neutral between his two options, his decision was still sufficiently rational to count as L-free, and free in the ordinary sense.¹³

So given that torn decisions are weakly and teleologically rational, and given that strong rationality isn't required for appropriate non-randomness or L-freedom, can I conclude now that torn decisions are sufficiently rational to be appropriately non-random and L-free? Well, one might object to my drawing this conclusion here on the grounds that there might be a fourth kind of rationality that is (a) required for L-freedom but (b) not possessed by torn decisions. But I find this hard to believe, for two related reasons.

The least important of the two reasons is that I just don't think much rationality at all is required for a decision to be L-free, or free in the ordinary sense. A decision can be irrational in all sorts of ways and still be L-free and ordinary-language free. We have no problem with the sentence, 'She chose of her own free will, but she chose irrationally.' On the other hand, the sentence 'She chose of her own free will, but she didn't author or control her choice' seems incoherent to us. And this suggests that rationality is just not as central to the notion of free will, or to the notion of L-freedom, as authorship and control are. In fact, I even have doubts about whether weak and teleological rationality are required. Here's a pair of *prima facie* arguments for thinking they might not be:

- (i) Suppose that Sally is deliberating about whether to take a job and comes to believe that her reasons for refusing the job outweigh those for taking it (and suppose that she has not overlooked any forgotten or subconscious reasons). But suppose that when she calls to turn the job down, she is suddenly overcome by a feeling that she ought to take it, and she acts on this impulse and decides on the spur of the moment to accept the job. *Prima facie*, it seems that this decision is not weakly rational; but this does not seem to be a reason to deny that it is L-free, or ordinary-language free.

(ii) Suppose that Floyd is reading a book, and he suddenly decides to go for a walk, simultaneously perhaps, with his acquiring a desire to go for a walk. *Prima facie*, it seems plausible to say that this decision is not teleologically rational, but again, this doesn't seem to give us any reason to deny that it is L-free, or ordinary-language free.

I will not pursue these *prima facie* arguments here, because in the present context, it doesn't matter whether weak or teleological rationality is required for L-freedom, because we've already seen that our torn decisions are weakly and teleologically rational. But considerations like these suggest that not much rationality is required for L-freedom, or ordinary-language freedom. It may be that weak and teleological rationality are required (i.e., that in order for a decision to be free, it has to be made for a reason and cannot flagrantly contradict the agent's consciously remembered reasons for choosing), but it seems unlikely that any more rationality than this is required. Again, a decision can be irrational in all sorts of ways and still be L-free and ordinary-language free.

The second and more important reason why I doubt that there is a fourth kind of rationality that is (a) required for L-freedom but (b) not possessed by torn decisions is this: when we reflect on actual torn decisions like Ralph's and Jane's, it seems intuitively that they just *are* rational enough to count as free. Indeed, both of these decisions seem *very* rational: in the first place, there doesn't seem to be much that Ralph or Jane could have done to have been more rational (would it have been more rational for Jane to have stared at her waiter until she believed that she had compelling reasons for one of her options?); and second, if we compare torn decisions to decisions that are fully determined by reasons, the only kind of rationality that the former seem to lack is strong rationality, which we've already seen is not required for L-freedom, or for ordinary-language free will.

I conclude, then, that ordinary torn decisions are appropriately rational, i.e., rational enough to fulfill the rationality requirement for appropriate non-randomness. And notice that I argued this point independently of any claims about whether our torn decisions are undetermined. Thus, the conclusion holds regardless of whether our torn decisions are undetermined, and so we have the result that *if our torn decisions are undetermined at the moment of choice in the way we're assuming here, then they are rational, in an appropriate sense of the term*. And this, of course, is just what I wanted to establish in this section.

Before going on, I want to respond to Dennett's claim that libertarians require indeterminacy prior to the formation of intention, for otherwise, deliberation will be rendered pointless. He says:

It would be insane to hope that after...deliberation had terminated with an assessment of the best available course of action, indeterminism would then intervene to flip the coin before action (1978, p. 51).

I have already hinted at how we can respond to this worry, for in arguing that Ralph's decision was teleologically rational, I argued that his reasons were relevant to his choice, and we can maintain that his deliberation was relevant to his choice in essentially the same way. We can say that Ralph's choice was made *in the light of* his deliberation and his reasons. Deliberation left him with the feeling that two of his options (moving to New York and staying in Mayberry) were better than all the others (e.g., moving to Outer Mongolia) and that he was torn between those two options. Given this, there could have been three appropriate responses to the deliberation: (a) just choose to move to New York; (b) just choose to stay in Mayberry; (c) just choose to deliberate some more. So by going with (a), Ralph didn't do anything that made his deliberation otiose (or that ignored his reasons). Moving to Mongolia would have been inappropriate, given his deliberation and reasons; but choosing to move to New York in a torn and perhaps undetermined-at-the-moment-of-choice sort of way was not inappropriate at all; indeed, given his deliberation and his reasons, it makes perfect sense. In short, even if Ralph's decision was undetermined at the moment of choice, and hence, not determined by his reasons or deliberation, it doesn't follow that his reasons or deliberation were irrelevant to the choice, and in fact they weren't, because they succeeded in (a) limiting his options to two, (b) providing reasons for both of these options, and (c) revealing that these two options were equally appealing.

Finally, I want to say a few words about whether my view is event-causal or noncausal. I have argued here that reasons are causally relevant to our torn decisions, because (i) they can be what causes a torn decision to occur at all, because a person can have good reasons to choose *now*, as opposed to remaining in a state of indecision; and (ii) they can cause us to choose from among a certain set of options, namely, our reasons-based tied-for-best options. But if our torn decisions are undetermined at the moment of choice in the way we're assuming here—i.e., if they are “wholly undetermined” in the sense specified in section 2.2—then which tied-for-best option is chosen is uncaused. So there is a noncausal element in the sort of libertarianism I'm describing. In the end, I think it's more natural to describe this view as event-causal than noncausal, because of (i) and (ii) above (and because, as we'll see in section 4, I think that when it comes to non-torn decisions, reasons can causally determine which option is chosen). Nonetheless, there is a noncausal element here: if torn decisions are undetermined at the

moment of choice in the way we're assuming, then they have properties that nothing causes them to have, e.g., properties like being-a-New-York-over-Mayberry-decision. (I suppose one might say that in the picture I'm constructing, *Ralph* caused his decision to have this property. But this would just be stipulative: if I added this claim to my libertarian view, I would be changing the view in only a terminological way; I would not be adding a substantive claim about the nature of reality.)

3.3. *Plural Authorship, Control, and Rationality*

So far, I have argued that if our torn decisions are undetermined at the moment of choice, then (a) we author and control them, and (b) they are sufficiently rational to count as L-free. The last thing I need to argue is that if a torn decision is undetermined at the moment of choice, then it satisfies the requirements of plural authorship, control, and rationality—that is, there is at least one other option that the person in question could have chosen such that if she had, then she still would have authored and controlled the decision, and it still would have been sufficiently rational to be L-free. This is easy to appreciate. Consider the case of *Ralph*. Given that his decision was undetermined at the moment of choice, if he had “just chosen” to stay in Mayberry, we would want to say the same thing about this decision that I have said here about his decision to move to New York: by the reasoning used in section 3.1, it would have been authored and controlled by *Ralph*, and by the reasoning used in section 3.2, it would have been sufficiently rational. The crucial point to notice here is that as long as we're dealing with torn decisions, plural authorship, control, and rationality are no harder to come by than one-way authorship, control, and rationality.

My conclusion, then, is that if any of our ordinary torn decisions are undetermined at the moment of choice, then they are appropriately non-random, because (a) we (plurally) author and control them, and (b) they are (plurally) sufficiently rational—in particular, they are (plurally) weakly and teleologically rational.

4. Non-Torn Decisions

I will not try to argue this point here, but I think that my thesis applies not just to torn decisions, but to various other kinds of decisions as well. For instance, I think my thesis applies to spontaneous, no-real-reason decisions like *Floyd's*, and I think it applies to *Buridan's-ass* decisions, i.e., decisions that are like torn decisions except that the reasons for the various tied-for-best options are the *same* reasons. (E.g., if *Bill* is waiting for an elevator, and two cars arrive at the same time, he might hesitate for a moment, and then consciously and purposefully choose to enter elevator 1 instead of elevator 2. In this case, *Bill* does have a reason for choosing to get on

elevator 1—namely, to get to the 7th floor, or wherever he's going—but this is also a reason for getting on elevator 2, and he doesn't have a reason for choosing elevator 1 over elevator 2.) For reasons very much like the reasons I have given here concerning torn decisions, I think that if decisions like *Floyd's* and *Bill's* are undetermined at the moment of choice, then they are appropriately non-random and, hence, L-free. But again, I will not try to argue this point here.

However, I would like to say a few words about cases in which the agent's reasons pick out a unique best option and perhaps even causally determine the agent's decision. I think that decisions like this can be free in the ordinary sense of the term, and I think that libertarians ought to allow that they can be L-free. We can allow for this very easily by saying something like the following: an *L-free person* is a person who is capable of making *type-1 decisions*, where a *type-1 decision* is defined as a decision that is (a) undetermined at the moment of choice and (b) appropriately non-random, where this involves at least (plural) authorship and control and perhaps also (plural) weak and teleological rationality; and an *L-free decision* is a decision made by an L-free person that is either a *type-1 decision* or a *type-2 decision*, where a *type-2 decision* is defined as a decision that is (a) determined by the agent's reasons for choosing and (b) appropriately non-random in the sense that it is (one-way) authored and controlled by the agent (and since such decisions are determined by the agent's reasons, they are also (one-way) weakly and strongly rational and (at least one-way) teleologically rational). Now, by saying this, I have not altered the definition of libertarianism; it is still just the thesis that human beings are L-free, and this is still defined as the ability to make *type-1 decisions*. I have not changed this, because I am not saying that *type-2 decisions* are *required* for L-freedom; on the view I have in mind, only *type-1 decisions* are required for L-freedom. My point is simply that libertarians can maintain that if we do make *type-1 decisions* and are, therefore, L-free, then our *type-2 decisions* can be considered L-free as well. Why not? There seems no reason not to say this, and saying it seems attractive because it fits with ordinary usage of the term ‘free will’.¹⁴

It is worth noting, too, that the distinction between *type-1* and *type-2* decisions is not precise. There is a spectrum of possible decision types, ranging from pure torn decisions on the one hand (i.e., cases where the agent's reasons are neutral between the various tied-for-best options) to pure non-torn decisions on the other (i.e., cases where the reasons pick out a unique best option); in between there are cases where the agent's reasons favor one or more of the live options and so the reasons-based probabilities of these options being chosen are higher than those of the other live options. This is not a problem. So long as some decisions that are purely torn (or at least reasonably close to purely torn) are undetermined at the moment of choice, libertarianism will be true, and we will be free to take all of the different kinds of decisions on this spectrum to be L-free, provided that

nothing external to the agent is causally responsible for the choice. E.g., if Sam's reasons favor option A over option B by a 70–30 margin, then in order for a choice of one of these options to be fully L-free, the moment-of-choice probabilities of A and B being chosen must be .7 and .3, respectively, given the complete state of the universe and the laws of nature; if something external to Sam determines which option is chosen, then the decision is not L-free, and if something external to Sam probabilistically causally influences the decision in the sense that it alters the moment-of-choice probabilities from .7 and .3 to something else, say .8 and .2, or .6 and .4, then the decision is not L-free to the extent that it is causally influenced by the external factor. (Of course, it is pretty artificial to hang precise numbers on the strengths of our reasons in this way; but this is really just a device for getting clear about what we ought to say about different kinds of cases.)

5. Indeterminism

We saw in section 3 that the question of whether libertarianism is true reduces to the question of whether any of our torn decisions are undetermined at the moment of choice. (More precisely, we can say that (a) the question of whether libertarianism is wholly true reduces to the question of whether any of our torn decisions are wholly undetermined, i.e., are such that at the moment of choice, which option is chosen is uncaused, so that the moment-of-choice probabilities of the various reasons-based tied-for-best options are at least roughly even, given the complete state of the world and all the laws of nature; and (b) the question of the extent to which libertarianism is true reduces to the question of the extent to which our torn decisions are causally undetermined at the moment of choice.) But I haven't argued that any of our torn decisions *are* undetermined at the moment of choice, and so for all that's been said here, it could be that humans are not L-free and, hence, that libertarianism is false. This is a straightforward empirical question: Are any of our torn decisions undetermined at the moment of choice? (Or more explicitly, (a) are any such decisions wholly undetermined in the above way?; and (b) if any such decisions are neither wholly undetermined nor strictly determined, to what extent are they undetermined?)

What I want to point out now is that these are *wide-open* questions. In particular, what's most important here is that, as of right now, we don't have any reason at all to disbelieve the claim that at least some of our torn decisions are wholly undetermined in the above way. The first point to note in this connection is that the most widely accepted interpretations of quantum mechanics take the world to be indeterministic. Now, there is one reputable interpretation that is deterministic—viz., Bohm's hidden variables interpretation—but as of right now, there is no good *evidence* for believing that it is true. Thus, for all we know right now, it *may* be that some quantum events are not strictly determined. Of course, it might also be

that they're all determined, and even if they're not, it may be that all of our torn decisions are determined. Or it may be that these decisions are *determined for all practical purposes*; e.g., it may be that the neural events that are our torn decisions are composed of large numbers of indeterministic quantum events that “cancel each other out”, so that for any torn decision, there is some specific option (e.g., moving to New York) such that prior circumstances together with causal laws make it overwhelmingly likely that the given option will be chosen. But again, as of right now, there is no good evidence for any such claim. In short, there is no good reason to believe that all of our torn decisions are, at the moment of choice, significantly causally influenced by factors external to us and our reasons. For all we know right now, it may be that (a) some of our torn decisions are not strictly determined—in particular, they would be huge conglomerates of quantum events, at least some of which are not strictly determined—and it may be that (b) some of these decisions are such that, at the moment of choice, which option is chosen is uncaused, so that the moment-of-choice probabilities of the various reasons-based tied-for-best options are at least roughly even, given the complete state of the world and all the laws of nature. As of right now, we have no good reason to doubt this, because we just don't know enough about quantum events, or neural events, or the relationship between the two, to have a well-informed view of whether (a) and (b) are true or not.

If this is correct, then the question of whether any of our torn decisions are undetermined (i.e., are such that, at the moment of choice, which option is chosen is uncaused, so that the moment-of-choice probabilities of the various tied-for-best options are at least roughly even, given the complete state of the world and the laws of nature) is a wide-open empirical question. But I have argued in this paper that the question of whether libertarianism is true reduces to this very question (and also to the more general (but still wide-open) empirical question of the extent to which our torn decisions are causally undetermined at the moment of choice). Thus, if I am right, then the question of whether libertarianism is true is a wide-open empirical question.¹⁵

Notes

¹ See Hobbes (1651), Hume (1748), and Hobart (1934).

² Two points: First, van Inwagen (1983) and the later Chisholm (1995) have developed event-causal libertarian views, but I do not include them on the list in the text, because, as far as I can tell, neither of these views involves a commitment to non-Valerianism (or to Valerianism). Second, some people, e.g., Clarke (2002), have taken Ekstrom's view to be Valerian; but I think this is because her terminology is confusing. She places the important indeterminism in *preference formation*, and on the most standard philosophical usage, this would suggest that it is located prior to decision making. But for Ekstrom (2000, p. 107), decision making is a *kind* of preference formation—it produces preferences concerning what we want to do—and it is

here that she places the important kind of indeterminism. So I think that Ekstrom is a non-Valerian libertarian, and in private correspondence she has endorsed this interpretation.

³ Libertarians needn't maintain that *all* our torn decisions are undetermined, because they're claiming only that we make *some* L-free decisions. (I suppose one might claim that the real question is whether *many* of our torn decisions are undetermined, since it would hardly vindicate libertarianism if it turned out that exactly one human decision was L-free; but I will assume that there is some degree of regularity here, so that if any of our torn decisions are undetermined, then many of them are.)

⁴ As long as the eliminativist allows that when humans have the experiences that we call decisions there are neural events occurring, I could define libertarianism, and run my argument, in terms of these neural events, regardless of whether they count as *decisions*.

⁵ The notion of a *torn decision* that I have defined in this section is similar to Kane's (1996) notion of a *self-forming action* (SFA). But there is a very important difference. SFAs are by definition undetermined, and so it is an open question whether there actually exist any such things. Torn decisions, on the other hand, are not defined as being undetermined; thus, it is clear that there do exist torn decisions—we know this by experience—but it is an open question whether any of our torn decisions are undetermined.

⁶ Kane first introduced conditions of this sort in his (1985), where he called them *duality* conditions. In his (1996), he called them plurality conditions.

⁷ There are important differences between the various formulations of this argument; e.g., van Inwagen used it to attack the notion of agent causation, whereas Clarke and O'Connor think the problem can be solved by invoking agent causation.

⁸ One might claim that torn decisions could be determined by the agent's subconscious reasons and still be authored and controlled by the agent, since subconscious reasons aren't extra-agential. I don't need to take a stand on this here, because I am trying to argue only that if our torn decisions are undetermined then they're appropriately non-random; I don't need the result that if they're determined then they're not appropriately non-random. If I were trying to motivate thesis (B) from section 2.2, then I would have to discuss this; but that will have to wait for another occasion.

⁹ The sort of control that Kane says is diminished by indeterminism is *antecedent determining control*—i.e., “the ability to determine or guarantee which of a set of options will occur *before* it occurs” (1999, p. 238); one might wonder how I could question Kane's claim that *this* kind of control is diminished by indeterminism. Let me say two things about this, one corresponding to point (i) in the text, and the other corresponding to point (ii). My first point is this: as long as we're talking about a torn decision, determinism wouldn't deliver antecedent determining control, because it wouldn't be the *agent* who did the prior determining, because before the choice, the agent is torn; thus, in the case of torn decisions, indeterminism doesn't diminish antecedent determining control because we wouldn't have this sort of control anyway. Moving on to my second point, since we don't have antecedent determining control in connection with torn decisions (whether they're undetermined or not), it might seem that point (ii) in the text is false—i.e., it might seem that we don't have as much control over torn decisions (whether they're undetermined or not) as we do over non-torn decisions in which our reasons determine our choices, because with torn decisions we don't have antecedent determining control. But I think this is a misconception. The Ralph₁/Ralph₂ argument suggests that a lack of antecedent determining control doesn't entail a lack of control, or even a decrease in the amount of control; just because Ralph₁ didn't determine his choice *in advance*, it doesn't follow that he didn't control the choice, and indeed, it doesn't follow that he had *less* control than he would have if he had determined it in advance. The question of whether an agent has antecedent determining control is a question about *when* an agent has control; it is not a question about *whether* an agent has control, or about *how much* control an agent has. If Ralph's decision was undetermined at the moment of choice, then he decided consciously, purposefully, and on his own, without anything making him choose as he did; thus, it seems to

me that (on the assumption that his choice was undetermined) Ralph₁ had *full* control—as much control as Ralph₂—despite the fact that he didn't have antecedent determining control. In short, if Ralph's decision was undetermined at the moment of choice, then what we should say is that he controlled it *at the moment of choice* but not beforehand.

¹⁰ On the standard usage, it is built into the definition of 'agent-causal view' that agent causation is not reducible to event causation. But it has always seemed to me an important feature of event-causal libertarianism that (if properly formulated) it can provide an intuitively pleasing conception of agent causation in event-causal terms. Ekstrom (2000, p. ix) and the later Chisholm (1995) have made similar points; they characterize themselves as providing event-causal accounts of agent causation. I should note, however, that their accounts are different from mine.

¹¹ If a person's reason set is inconsistent, then on these definitions, all of her decisions will be strongly rational but not weakly rational. In connection with strong rationality, we can fix this problem by demanding that either (a) the reason set is consistent and entails that the given decision is best or (b) the reason set is inconsistent but for all of the “reasonable ways of doing away with the inconsistency”, the resulting reason set still entails that the given decision is best. And we could make an analogous adjustment to the definition of 'weak rationality'. Of course, much would need to be said about the notion of a “reasonable way of doing away with an inconsistency”, but I haven't the space to go into this here.

¹² There is a slight problem with these definitions of overall-strong and overall-weak rationality that is analogous to the problem with weak and strong rationality discussed in note 11. But the problem here can be solved in a way that is essentially analogous to the way in which I solve the problem there.

¹³ A similar argument applies to decisions that fly in the face of reason sets that include reasons which are conscious in the sense of being accessible to the agent, but which were simply *forgotten* during deliberation and choice.

¹⁴ Kane (1996) also allows that type-2 decisions can be L-free; his view is a bit different from mine here, but I do not have the space to go into this.

¹⁵ I would like to thank John Martin Fischer, Carl Ginet, Robert Jones, Robert Kane, Storrs McCall, Timothy O'Connor, and two anonymous referees for very helpful comments on earlier versions of this paper.

References

- Bernstein, M. (1995) “Kanean Libertarianism,” *Southwest Philosophy Review* 11:151–57.
- Berofsky, B. (2000) “Ultimate Responsibility in a Determined World,” *Philosophy and Phenomenological Research* 60:135–40.
- Chisholm, R. (1995) “Agents, Causes, and Events: The Problem of Free Will,” in T. O'Connor, ed., *Agents, Causes, and Events*, New York: Oxford University Press, pp. 95–100.
- Clarke, R. (1993) “Toward a Credible Agent-Causal Account of Free Will,” *Noûs* 27:191–203.
- . (1995) “Indeterminism and Control,” *American Philosophical Quarterly* 32:125–38.
- . (2002) “Libertarian Views: Critical Survey of Noncausal and Event-Causal Accounts of Free Agency,” in Kane (2002).
- Dennett, D. (1978) “On Giving Libertarians What They Say They Want,” in *Brainstorms*, Cambridge, MA: MIT Press.
- Double, R. (1991) *The Non-Reality of Free Will*, New York: Oxford University Press.
- Ekstrom, L. W. (2000) *Free Will: A Philosophical Study*, Boulder, CO: Westview Press.
- Fischer, J. M. (1999) “Recent Work on Moral Responsibility,” *Ethics* 110:93–139.
- Ginet, C. (1990) *On Action*, Cambridge: Cambridge University Press.
- . (2002) “Reasons Explanations of Action: Causalist versus Noncausal Accounts,” in Kane (2002).
- Goetz, S. (1997) “Libertarian Choice,” *Faith and Philosophy* 14:195–211.

- Haji, I. (1999) "Indeterminism and Frankfurt-style Examples," *Philosophical Explorations* 2:42-58.
- Hobart, R. E. (1934) "Free Will as Involving Determinism and Inconceivable Without It," *Mind* 43:1-27.
- Hobbes, T. (1651) *Leviathan*. Reprinted 1962, New York: Collier Books.
- Hume, D. (1748) *An Inquiry Concerning Human Understanding*. Reprinted 1955, Indianapolis, IN: Bobbs-Merrill.
- Kane, R. (1985) *Free Will and Values*, Albany, NY: State University of New York Press.
- . (1996) *The Significance of Free Will*, New York: Oxford University Press.
- . (1999) "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism," *Journal of Philosophy* 96:217-40.
- . ed. (2002) *The Oxford Handbook of Free Will*, New York: Oxford University Press.
- McCall, S. (1994) *A Model of the Universe*, Oxford: Oxford University Press.
- McCann, H. (1998) *The Works of Agency: On Human Action, Will and Freedom*, Ithaca, NY: Cornell University Press.
- Mele, A. (1995) *Autonomous Agents*, New York: Oxford University Press.
- . (1999a) "Kane, Luck and the Significance of Free Will," *Philosophical Explorations* 2:96-104.
- . (1999b) "Ultimate Responsibility and Dumb Luck," *Social Philosophy and Policy* 16:274-93.
- Nozick, R. (1981) *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- O'Connor, T. (2000) *Persons and Causes*, New York: Oxford University Press.
- Strawson, G. (2000) "The Unhelpfulness of Indeterminism," *Philosophy and Phenomenological Research* 60:149-55.
- van Inwagen, P. (1983) *An Essay on Free Will*, Oxford: Oxford University Press.
- . (2002) "Free Will Remains a Mystery," in Kane (2002).
- Waller, B. (1988) "Free Will Gone Out of Control: A Review of R. Kane's *Free Will and Values*," *Behaviorism* 16:149-62.