

The Two-Stage Luck Objection

SETH SHABO

University of Delaware

1. Introduction

Libertarians believe that free will exists and that its existence is incompatible with determinism. The leading challenge to libertarianism is the Luck Objection. While the challenge has taken various forms, the common theme is that libertarians cannot explain how causally undetermined actions could be exercises of free will, as opposed to mere matters of luck or chance for which no one is morally responsible.¹ Failing such an explanation, these theorists may need to choose between their belief in free will and their commitment to incompatibilism, or else conclude that free will is a real but mysterious feature of human agency (van Inwagen 2000).

In recent years, discussions of the Luck Objection have reached an impasse of sorts. On the one hand, existing responses to the objection have failed to satisfy libertarianism's many critics. On the other hand, a growing number of libertarians seem unimpressed by existing formulations of the objection.² My aim here is to break this impasse by sharpening the objection. To this end, I start by dividing it into two stages. The first and more familiar stage seeks to show that supposed exercises of free will are importantly similar to other causally undetermined outcomes for which no one is morally responsible. Just as which way a remote particle swerves, or which number a random-number generator generates, is a "mere matter of luck or chance" for which no one is morally responsible, the idea is, so supposed exercises of free will are mere matters of luck or chance as well. The second stage, which usually goes unstated, aims to show that we aren't morally responsible for supposed exercises of free will *because* they resemble these other outcomes.³

The step from the first stage to the second is so small that it can easily escape notice. For, it might be thought, if defenders of the Luck Objection can show that supposed exercises of free will will resemble these other outcomes—outcomes for which we aren't morally responsible—in being mere matters of luck or chance, they will have shown that we aren't morally responsible for supposed exercises of free will. If so, the argument's second stage is otiose.

As we shall see, this second stage, far from being otiose, is a valuable resource for defenders of the Luck Objection. This is in part because it eases the burden on the argument's first stage. For suppose it can be shown in the second stage that we aren't morally responsible for the relevant actions if they resemble swerving particles and the like in being mere matters of luck or chance—or, as I prefer to say, *truly random outcomes*. To say that an outcome is truly random is to say, roughly, that both its

occurrence and its nonoccurrence have an objective probability substantially greater than zero and substantially less than one, and that there is no further explanation, beyond the probability distribution and everything that accounts for it, for why things turn out as they actually do. In that case, the first stage need only establish that such actions *are* truly random outcomes; it need not establish the contentious further claim that we lack suitable “control” over these actions. This further claim is contentious not least because the operative notion of control remains elusive.⁴ Unlike standard versions of the Luck Objection, the two-stage approach eschews this further claim by arguing, first, that the relevant actions meet the conditions for being truly random outcomes, and, second, that we aren’t morally responsible for actions that meet these conditions.

Since I have argued elsewhere that supposed exercises of free will are truly random outcomes,⁵ the emphasis here will be on the second stage. To show that we aren’t morally responsible for the relevant actions if they are truly random outcomes, I present the No Guidance Argument, which can be summarized this way:

When choosing between two or more alternative courses of action, we are morally responsible for choosing as we do only if we can reasonably be expected to guide our choice by applicable normative standards—that is, to comply with such standards at least in part because we take them to apply to us. And just as we cannot reasonably be expected to guide truly random nonactions by such standards, so we cannot reasonably be expected to guide supposed exercises of free will by such standards if they are truly random outcomes, notwithstanding that they are actions and not mere happenings.

In short, the relevant actions, qua truly random outcomes, fail to meet a guidance requirement for morally responsible agency.

To be sure, some readers will think it obvious that we aren’t morally responsible for truly random outcomes as I define them, even as others demand to know, “Why not?” Whatever else, both groups should agree that we are better off with an explicit answer to this question. By providing such an answer, the No Guidance Argument fills an outstanding gap in the literature.

In section 2 below, I set out the first stage of the Two-stage Luck Objection. In section 3, I consider two possible libertarian responses. In section 4, I present the No Guidance Argument. In section 5, after connecting the two stages, I return to the libertarian responses from section 2.

2. Stage One—Free Will and Truly Random Outcomes

The Luck Objection’s first stage seeks to show that supposed exercises of free will share with swerving particles and randomly generated numbers the features that lead us to describe the latter as “mere matters of luck or chance.” In this section, I argue that supposed exercises of free will possess these features, despite being actions and not mere happenings. I shall term occurrences that possess these features ‘truly random outcomes’. An occurrence is a truly random outcome, as I shall use this expression, if and only if:

- (a) immediately before it occurs, an alternative occurrence is consistent with the total state of the world at that time, together with the laws of nature;
- (b) each possible occurrence has a definite, objective, “ground-floor” probability greater than zero and less than one;
- (c) neither occurrence is anomalous or exceedingly unlikely in the circumstances;
- (d) beyond the probability distribution and everything that accounts for it, there is no further explanation for which possible occurrence actually ensues.⁶

If, for example, a machine’s going into one of two possible internal states is a truly random outcome, there will be a determinate answer to the question, “How likely was each of the two possible states to ensue?” (say, 0.45 versus 0.55), but there will be no further explanation for why, given this probability distribution and everything that accounts for it, it goes into whichever state it actually does. Once we have cited the probability distribution and everything that accounts for it, explanation is at an end; there is “nothing more to say” about why things take the more specific turn they do.⁷

To further illustrate this last point, suppose that Christine activates the randomizing machine, thereby causing it to go into the one state or the other, and that it ends up going into the first state. We can say that it’s because Christine activates the machine that a random outcome ensues then, so that the machine goes into the one state or the other. We can even say that it’s because she activates the machine that it goes into the *first* state. With this last claim, however, we must be clear that we’re saying only this: the machine’s going into the first state is the outcome of the randomized process that Christine initiates, a process of which this state is one of two possible outcomes. What we cannot say is that things take the particular turn they do—with the randomized process issuing specifically in the first state—because she activates the machine. For activating the machine leaves it entirely open (subject only to the probability distribution) which of the two possible states actually ensues.

Put another way, we can cite the machine’s activation to explain why the question, “Which of the two states does the machine go into?” has an answer, whatever the answer is. But we cannot thereby explain why the answer is specifically “the first state,” since activating the machine does nothing to influence which way this question comes to be answered. If there were some further explanation for why the first state actually ensues, this explanation would go beyond the considerations already factored into the probability distribution; and so, by (d) above, the outcome wouldn’t be truly random after all.

Why believe that supposed exercises of free will would be truly random outcomes? We might summarize the argument this way:

Supposed exercises of free will are causally undetermined choices. Though actions, these choices resemble random nonactions, such as the machine’s going into the first state, in respects (a)–(c) above. And if they resemble random nonactions in (a)–(c), they’re bound to resemble them in (d) as well. Thus supposed exercises of free will are truly random outcomes.

In short, libertarians can neither distinguish the relevant choices from truly random outcomes in (a)–(c) nor distinguish them in (d) without distinguishing them in one of (a)–(c).

It should be noted that this statement of the argument doesn't address agent-causal libertarians. These theorists believe that agents *qua* agents (or substances) make a distinctive causal contribution to their free actions, a contribution "over and above" that of their states. These theorists can deny that the relevant choices meet (d) without denying that they meet (a)–(c). For if Christine, *qua* agent, causes her choice, there *will* be something more to say about why that choice ensues, namely that Christine, *qua* agent, causes it. While I believe that the argument from this section can be extended to agent-causal accounts, I won't pursue this argument here.⁸

Returning to the scenario above, suppose for discussion's sake that Christine's choice to activate the machine is an exercise of free will. We might imagine that activating it indicates her agreement to bet a considerable sum of money on the outcome of the randomized process she initiates, and that she is torn between her self-interested reasons for accepting the wager (she stands to gain considerably if things go her way) and her moral reasons for declining it (placing the bet would be a serious breach of her fiduciary responsibilities). Many libertarians will say that such a choice, *qua* exercise of free will, meets (a)–(c): each choice is consistent with the immediate past and the laws of nature; each has a definite, "ground-floor," objective probability significantly greater than zero and significantly less than one; and neither would be anomalous or exceedingly improbable given Christine's torn state.

Now accepting both (a) and (c) doesn't automatically commit libertarians to (b). A libertarian could conceivably deny, for example, that either choice would be anomalous or exceedingly improbable without accepting that there is a definite answer to the question, "How likely is each possible choice?" (See Ekstrom 2011, p. 376, and Buchak 2013.) Unlike some influential versions of the Luck Objection, such as the Rollback Argument (van Inwagen 2000), the Two-stage Luck Objection doesn't depend on such definite probability assignments. For what matters is not that the objective probabilities have definite values, such 0.6 versus 0.4, but that there is "nothing more to say" about why the agent chooses as she actually does, once all of her causally relevant mental states have been factored into the probability distribution, however we cash out that distribution (see note 6). For simplicity, however, I shall continue to suppose that the alternatives have definite probabilities.

This brings us to the argument's second premise: that the relevant choices meet (d) if they meet (a)–(c). That is, if Christine's choice to accept the wager meets (a)–(c), there will be nothing more to say, beyond citing the probability distribution and everything that accounts for it, about why she actually makes this choice in particular. Why believe this?

Recall that when explicating the notion of a truly random outcome, we saw that there couldn't be any further explanation (beyond the probability distribution and everything that accounts for it) for why Christine chooses specifically to activate the machine; by definition, any such further explanation would disqualify the choice from being truly random. The goal now is not to explicate this notion, but to argue that the relevant choices would indeed *be* truly random outcomes—in other

words, that they would meet (d) given that they meet (a)–(c). Nevertheless, the random-by-stipulation choice will serve as a useful foil. In effect, the argument’s second premise challenges libertarians to locate the crucial difference between supposed exercises of free will and such truly random choices, a difference that provides “something more to say”—an explanation that goes beyond the probability distribution and everything that accounts for it—about why someone makes the particular choice she does in a supposed exercise of free will. Appeals to agent-causation aside, I contend that libertarians won’t find such a difference.

To see why, consider the following scenario:

Christine is torn between her self-interested reasons for activating the machine and thereby accepting the wager, and her moral reasons for declining the wager by refraining from activating it. Immediately before she chooses, her inclination to accept is somewhat stronger than her competing motivations, with a “ground-floor,” objective probability distribution of 0.55 versus 0.45. As things turn out, she chooses to accept it—which she does by activating the machine—for her self-interested reasons. Had she chosen to decline the wager instead, her choice would have been based on her moral reasons for declining it. As to why she actually chooses as she does, there is simply nothing more to say once we have cited the probability distribution and all the mental states that account for it.

Since there is nothing more to say about why Christine ends up choosing to activate the machine in this scenario, her choice is a truly random outcome, just as whichever state the machine goes into is a truly random outcome.

To show that supposed exercises of free will are not truly random outcomes, despite resembling Christine’s choice in (a)–(c), libertarians must posit a further explanatory consideration—one that goes beyond the probability distribution and everything that accounts for it—for why the supposedly free agent makes the particular choice he does, given that distribution. And, I submit, libertarians will be hard pressed to find such a further explanatory consideration. For, having already factored the agent’s causally relevant mental states into the probability distribution, they cannot appeal to any combination of these states to explain why, given this distribution, she actually chooses as she does. Thus, if libertarians’ explanatory resources are confined to agents’ states, they will be unable to distinguish supposed exercises of free will from choices like Christine’s.

Note that my conclusion is not that we would lack suitable control over which choice we make, or that such choices would not be exercises of free will.⁹ The conclusion is, rather, that the relevant choices—the ones that libertarians take to be exercises of free will—are truly random outcomes, whatever else they may be.

To be sure, some libertarians will say that this conclusion, if true, effectively rules out free will. As we shall see, however, other libertarians will dispute this assessment. In any case, the important point for now is that ‘Such choices are truly random outcomes’ and ‘Such choices aren’t exercises of free will’ express distinct propositions, and the above argument purports to establish only the first. If libertarians wish to resist this first stage of the Two-stage Luck Objection, they must deny that choices like Christine’s resemble the machine’s outcome in one of (a)–(d).

3. Libertarian Responses to the First Stage

Let us now consider two possible libertarian responses to this argument. The first accepts that these choices are truly random outcomes, while maintaining that they are *also* exercises of free will. For, the idea is, an exercise of free will is nothing more than a truly random outcome that someone does for her reasons. In other words, what separates exercises of free will from truly random nonactions is only that the former, being choices, are attributable to us as our actions, things we do for our reasons. When Christine chooses to activate the machine, she thereby manifests her ability to make this choice for her reasons, when she could have manifested her ability to choose otherwise, for her countervailing reasons.¹⁰ If this is what free will consists in, libertarians can maintain that Christine is morally responsible in the above scenario, without denying that her choice is a truly random outcome. I take up this response in section 4, where I argue “directly”—that is, without reference to free will—that we aren’t morally responsible for such choices if they are truly random outcomes.

The second libertarian response maintains that exercises of free will are crucially different from Christine’s choice after all. For, given a suitably detailed account of the causally relevant mental states in a genuine exercise of free will, we can plausibly see the ensuing choice as manifestation of the agent’s control over which choice ensues, control that Christine lacks. This control, or the fact that the agent exercises it as he does, can then be cited to account for why he makes that particular choice. In addition to the fact that he acts for his reasons plus the facts that determine the probability distribution, these libertarians will say, there is this further “fact of specific control” (as we might call it). If libertarians can maintain that this fact is indeed something “over and above” performing a causally undetermined action for one’s reasons, they can deny that such choices meet (d), thereby distinguishing exercises of free will from choices like Christine’s.

For libertarians who favor this second response, the worry is not that this supposed further “fact of specific control” is unlikely to provide an illuminating explanation of why someone makes the particular choice she does. The worry is, rather, that there seems to be no room for such a further fact at all. For event-causal libertarians can appeal only to agents’ states, and any combination of states they appeal to will have exhausted its causal influence over the outcome in determining the probability distribution. It’s thus hard to see how a story that features only an agent’s states could provide something more to say about why she makes the specific choice she does, given the probability distribution and everything that accounts for it. Put another way, it looks like we can import the details of any such event-causal account into Christine’s scenario, consistent with her choice’s remaining a truly random outcome of her prior mental states. For this reason, such a further “fact of specific control” seems unlikely to pan out, and so the second response isn’t a viable alternative to the first.

One advantage of the two-stage approach is that it forces event-causal libertarians to squarely confront this question: “When someone’s choice is an exercise of

free will, is there some further fact—a fact beyond the ones that determine the probability distribution, plus the agent’s mental states’ appropriately causing her choice—that helps to explain which choice actually ensues, given the probability distribution? In pressing this question, we are pressing them to clearly commit to either the first or the second response.¹¹

Again, seeing no room for such a further fact, I don’t see how event-causal libertarians can differentiate supposed exercises of free will from truly random outcomes. If they can’t, the first stage of the Two-stage Luck Objection undercuts the second libertarian response. The second stage of the objection takes aim at the first, “reductive” (or “minimalist”) libertarian response. If supposed exercises of free will are nothing but truly random outcomes, I contend, libertarians cannot plausibly maintain that we are morally responsible for them, just as they cannot plausibly maintain that Christine is morally responsible for the randomizing machine’s going specifically into the first state when she activates it.

4. Stage Two—Random Outcomes and Normative Guidance

If the relevant choices are truly random outcomes, like the machine’s going into the first state, some will find it obvious that we aren’t morally responsible for them, even as others reply, “Why not?” Those in the second group will include libertarians who believe that exercises of free will are nothing over and above truly random actions that we perform for our reasons (see Balaguer 2010, especially pp. 107–11, and Franklin 2012 and forthcoming), along with some contemporary compatibilists who wish to make moral responsibility safe for both determinism and indeterminism (see especially Mele 2006, ch.5, and Fischer 2012, ch. 6, and 2014). And it will include many others who find the claim intuitively plausible but who still need convincing.¹²

To answer all those in the second group, we need an argument that picks up where the one from the previous section left off. According to the No Guidance Argument, such choices don’t meet a plausible guidance requirement for morally responsible agency. I shall now set out this argument.

4.1. The No Guidance Argument’s First Three Premises

Returning to Christine’s choice to activate the machine, recall that the issue is whether Christine is directly morally responsible for her choice, not whether she is indirectly or derivatively responsible for it—say, by virtue of being morally responsible for coming to face this choice in the first place. I will have direct moral responsibility in mind in what follows.

In keeping with the argument from the previous section, assume that Christine’s choice to activate the machine is a truly random outcome of her prior mental states. The No Guidance Argument’s first premise is that

- (N1) Christine is morally responsible for her choice only if she is morally responsible for choosing as she does.

Given that Christine actually chooses to activate the machine, N1 implies that she isn't morally responsible for her choice unless she is morally responsible for choosing specifically to activate the machine, or for the fact that she so chooses. To say that Christine is morally responsible for choosing as she does is to say that she is morally responsible for how things actually turn out with respect to which choice she makes. In other words, it is to say that she is morally responsible for *which* of the possible choices actually ensues, or for how the question, "Which possible choice actually ensues?" comes to be answered. If she is not morally responsible for which choice actually ensues, then she is not morally responsible for the ensuing choice, whichever choice she makes. More generally, when we speak of moral responsibility for choices, we are speaking of moral responsibility for making the specific choice one does: to lie or to tell the truth, to vote one way or another, to honor a commitment or to renege on it, and so on.

If I'm right, this premise, which is rarely made explicit, is crucial to understanding the problem about moral responsibility for random choices. For the issue is precisely whether agents are morally responsible for how these choices turn out, or for how the question, "Which choice will ensue?" actually comes to be answered. If Christine isn't morally responsible for which way she ends up choosing, she isn't morally responsible for choosing specifically to activate the machine; and so, by N1, she isn't (directly) morally responsible for this choice simpliciter.

I take the argument's next two premises to be uncontroversial as well. According to the second premise,

- (N2) Christine is morally responsible for choosing as she does only if she is subject to normative expectations concerning which choice she makes.

When Christine chooses to accept the wager, she chooses to do something she knows to be morally wrong. If she is morally responsible for this choice, there is at least some extent to which she is morally blameworthy for it. And if she is morally blameworthy for so choosing, she must be subject to normative standards or expectations, standards or expectations that apply to her as the maker of that choice, and that she violates by choosing to accept the wager. If Christine isn't subject to such standards and expectations concerning which choice she makes, there are no grounds for judging her to be morally blameworthy for choosing as she does. Thus, she is morally responsible for her choice only if she is subject to normative expectations concerning which choice she makes.

Next,

- (N3) Christine is subject to normative expectations concerning which choice she makes only if she can reasonably be held to such expectations.

When we speak of holding someone to a normative expectation, we usually have one of two things in mind. First, we may mean that we are disposed to various appraisals and reactions—such as judging the person to be morally blameworthy or at fault, expressing blame and reproach or criticism, pressing him to account for his conduct, and so on—in the event that he violates this expectation on some occasion

or other. Second, we may be referring to an actually occurring response of this sort, one that is occasioned by a specific breach of our expectations. I am interested here in the second, more specific use of ‘holding someone to an expectation’. According to N3, Christine is subject to normative expectations concerning which choice she makes only if such appraisals and reactions are reasonable once she has actually chosen contrary to them. Put another way, Christine is subject to these expectations concerning her choice only if she is eligible for moral blame and the like because she has chosen contrary to them.

Of course, it’s a further question whether expressions of blame and reproach are justified all things considered. But if Christine isn’t eligible for such reactions in the first place, this further question doesn’t arise. On N3, then, Christine isn’t subject to the relevant normative expectations unless choosing contrary to them is a legitimate basis for such appraisals and reactions. Even if such reactions aren’t justified all things considered, they won’t be unreasonable or misplaced as they would be if her choice complied with our expectations or if those expectations didn’t apply to her (say, because she suffered from a psychological disorder or impairment). If her choice to accept the wager, a choice she knows to be morally wrong, doesn’t at least make her eligible for such reactions, she isn’t subject to normative expectations concerning which choice she makes.

Taken together, the argument’s first three premises imply that Christine is morally responsible for her choice to accept the wager only if she can reasonably be held to the normative expectation that she choose to decline the wager, so that she is eligible for moral blame on the basis of choosing contrary to this expectation.

4.2. *The No Guidance Argument’s Last Three Premises*

The argument’s fourth premise connects the reasonableness of holding Christine to this expectation concerning her choice with the reasonableness of a second normative expectation: the expectation that she guide her choice by the first expectation. I shall dub them ‘the First Expectation’ and ‘the Second Expectation’, respectively. Thus we have

The First Expectation: that Christine choose to decline the wager

The Second Expectation: that Christine guide her choice by the First Expectation

What do I mean by ‘Christine guides her choice by the First Expectation?’ As I shall use this expression, to say that Christine guides her choice by the First Expectation is to say that she chooses to decline the wager at least in part because she believes that she ought to decline it, where the ‘because’ is causal. In other words, it is to say that a belief about what she ought to do—namely, decline the wager—is among the causes of her ensuing choice to decline the wager. Thus, if Christine guides her choice by the First Expectation, two things are true of her:

- (1) She chooses to decline the wager, in accordance with the First Expectation.
- (2) Her belief that she ought to decline the wager¹³ is among the causes of this choice.

To say that the Second Expectation is reasonable, then, is to say that Christine can reasonably be expected to choose to decline the wager, in accordance with the First Expectation, at least in part because she believes that she should decline it, so that this belief is among the causes of her so choosing.

According to the fourth premise,

- (N4) Christine can reasonably be held to the First Expectation only if the Second Expectation is reasonable in her circumstances as well.

That is, we can reasonably hold Christine to the expectation that she choose to decline the wager only if we can also reasonably expect her to guide her choice by the expectation that she decline the wager—that is, to choose to decline it at least in part because she believes that she should decline it.

Note that the requirement expressed in N4 is not that Christine must successfully guide her choice by the expectation that she decline the wager, or even that she must actually believe that she should decline it, but only that these things *can reasonably be expected of her*. If someone assigns insufficient weight to a moral consideration and ends up flouting it, it may be perfectly reasonable to say that he should have assigned it greater weight and guided his choice accordingly. The same is true if someone fails to apply a relevant moral consideration to her situation: the expectation that she apply it may still be in order. In short, the reasonableness of expecting such guidance is no guarantee that it will actually occur. N4 connects the reasonableness of holding Christine to the First Expectation (that she choose to decline the wager) with the *reasonableness* of the Second Expectation (that the First Expectation guide her choice). Just as the First Expectation can be reasonable even if Christine doesn't comply with it, so the Second Expectation can be reasonable even if she doesn't actually guide her choice by the First Expectation.

Again, had Christine guided her choice by the First Expectation, she would have chosen to decline the wager at least in part because she believed that she ought to decline it. When Christine actually chooses to accept the wager, we hold her to the First Expectation by morally blaming or reproaching her, or at least by viewing her as eligible for such reactions, on the basis of her choice. On N4, holding her to this expectation isn't reasonable unless it's also reasonable to have expected her to guide her choice by it. If we believe that the First Expectation should have guided Christine's choice, so that we endorse the Second Expectation, we believe both that she should have seen that declining the wager was the right thing to do, and that her seeing this should have led her to choose to decline it. In other words, we believe that she should have chosen to decline the wager at least in part because she believed that she should decline it.

To be sure, there are important questions about when we can and cannot reasonably be expected to guide ourselves by normative expectations concerning how we choose and act. But we need not delve into these questions to support N4. For if we are genuinely convinced that Christine could not have reasonably been expected in the circumstances—circumstances in which both choices are possible—to guide her choice by the relevant normative expectations, so that she chooses in accordance

with those expectations at least in part because she believes that she should decline the wager, we should conclude that she isn't eligible for moral blame on the basis of choosing contrary to those expectations. Whatever else, she is not directly morally blameworthy on the basis of so choosing.

I have argued that Christine is subject to the First Expectation only if she can reasonably be held to it, and that she can reasonably be held to it, in turn, only if the Second Expectation is reasonable in her circumstances as well. It's important to see that these two normative expectations play different roles in assessing Christine's moral responsibility. In asking whether Christine is morally responsible for her choice, I have argued, we must ask whether she can reasonably be *held* to the First Expectation, the expectation that she choose to decline the wager. In holding her to this expectation, we view her as eligible for moral blame and the like because she has chosen contrary to it. We do not thereby *hold* her to the Second Expectation, the expectation that she guide her choice by the First Expectation. If she chooses to decline the wager, she will have satisfied our expectation concerning which choice she makes, regardless of whether she so chooses at least in part because she believes that she should decline the wager. Instead of holding Christine to the Second Expectation, then, we use it to assess the reasonableness of holding her to the First Expectation. If expecting Christine to guide her choice by our normative expectations would be unreasonable in her circumstances, the argument goes, she isn't eligible for moral blame on the basis of how she chooses, and so she cannot reasonably be held to those expectations in the circumstances.

Underlying N4 is a plausible principle about when we can reasonably hold people, in the form of morally blaming them and the like, to expectations they have violated. Consider someone who faces a choice between two alternative courses of action, where she will be free from external interference in adopting either course, and where she will exercise her ordinary capacity for self-governance (whatever exactly this amounts to) in selecting whichever course she does. Call someone who makes such a choice a 'self-governing agent'. We might put the principle this way:

(GP) We cannot reasonably hold a self-governing agent to our normative expectations concerning how she chooses and acts on some occasion unless we can reasonably expect her to guide her choice and action by those expectations on that occasion.

GP proposes an answer to this question: "Can we reasonably hold a self-governing agent to our expectations concerning which course she takes on some occasion, viewing her as eligible for moral blame and the like if she proceeds contrary to them?"¹⁴ The proposed answer is: "Not unless we could have reasonably expected her on that occasion to guide herself by the expectation to which we would be holding her—that is, to choose and act in accordance with that expectation at least in part because she believes that she should so choose and act."¹⁵

As to why we should accept GP, consider the situation of a *morally motivated* self-governing agent, who wants nothing more than to do what she is morally required to do. Apart from applying the relevant moral considerations to her

situation and thereby guiding herself towards compliance, such an agent has nothing to “steer by”: she has no means of bringing it about that she complies, despite her strong motivation to do so. If, for example, correctly applying the relevant considerations to her situation would require information that she blamelessly lacks, it isn’t reasonable to hold her to our expectation if she fails to comply. And, the thought continues, if a morally motivated self-governing agent cannot reasonably be held to our expectations concerning which choice she makes, given that she cannot reasonably be expected to guide her choice by the relevant moral considerations, the same will be true of self-governing agents more generally. We cannot reasonably hold such agents to expectations concerning how they choose and act on some occasion in full recognition that they cannot reasonably be expected on that occasion to guide themselves by the expectation to which we would be holding them.

By assumption, Christine is a self-governing agent: which course she actually takes comes down to how she exercises her powers of self-governance, and she will be free from external interference either way. If N4 is correct, it isn’t reasonable to hold her to our expectation concerning which choice she makes unless it’s also reasonable to expect her to guide her choice by the First Expectation—that is, to choose to decline the wager at least in part because she believes that she should decline it. If libertarians wish to reject this premise, it’s up to them to say why.¹⁶ But perhaps libertarians won’t dispute N4. Perhaps they will instead maintain that Christine can be expected to guide her choice by the First Expectation, notwithstanding that her choice is a truly random outcome of her prior mental states. The remainder of the No Guidance Argument seeks to rule out this possibility.

As we have seen, the guidance requirement expressed in N4 has a causal component: it must be reasonable to expect Christine to comply with the First Expectation at least in part because she believes that she should decline the wager. Now in the counterfactual event that Christine chooses to decline the wager, she will indeed have guided her choice by the First Expectation: she will have so chosen at least in part because of this belief. Thus it is possible that Christine will end up guiding her choice in this manner. But what is possible is not the same as what can reasonably be expected. Can such guidance reasonably be expected of Christine in her torn state?

Let us abbreviate ‘Christine’s belief that she should decline the wager’ as ‘the Belief’. According to the argument’s fifth premise,

- (N5) In her torn state, Christine can reasonably be expected to guide her choice by the First Expectation only if the Belief could causally influence the outcome of that state.

By assumption, Christine has the Belief, convinced as she is that accepting the wager would be morally wrong. According to N5, the Second Expectation isn’t reasonable in her torn state unless the Belief could causally influence which choice actually succeeds this state. Put another way, if the Belief could not causally influence the outcome in her actual, torn circumstances, she cannot reasonably be expected in

these circumstances to guide the ensuing choice by the First Expectation—that is, to choose to decline the wager at least in part because of the Belief. We cannot reasonably expect such guidance even though it may actually end up occurring.

To see why, suppose for discussion's sake that the Belief, having exhausted its causal influence, cannot causally influence the outcome of her torn state. How exactly is she supposed to bring it about that it guides the outcome of this state? What means does she have at her disposal for achieving this guidance? Whatever else, the Belief does not give her such a means. For whether this guidance occurs depends entirely on whether the Belief proves to be causally efficacious, and whether the Belief proves to be causally efficacious, in turn, depends entirely on which choice actually ensues. If a choice to decline the wager ensues, the Belief will count as causally efficacious; but, importantly, it will be the occurrence of this choice that *makes* it causally efficacious, not the other way around. After all, with Christine's competing mental states having exhausted their causal influence, the question, "Which set of mental states ends up probabilistically causing the ensuing choice?" reduces to the question, "Which choice actually ensues?" And, the thought continues, if the efficacy of the Belief depends entirely on which choice actually ensues, the Belief does not give her a means by which to bring it about that the ensuing choice complies with the First Expectation. Thus, if the Belief cannot causally influence the outcome of Christine's torn state, it does not give her a means of bringing it about that she guides her torn choice by the First Expectation. It does not give her such a means because whether or not it proves to be causally efficacious is entirely dependent on whether the one choice or the other actually ensues.¹⁷

Implicit in the preceding paragraph was a distinction between two states of affairs:

- (S1) The Belief is causally efficacious—that is, Christine chooses to decline the wager at least in part because of the Belief.
- (S2) The Belief causally influences the outcome of her torn state—that is, Christine chooses to decline the wager at least in part because of this causal influence.

While both S1 and S2 imply that the First Expectation guides Christine's choice, S2 is stronger than S1. For if the Belief causally influences the outcome of her torn state, as opposed to merely being among the causes of the ensuing choice, we could cite its presence to explain not only why Christine chooses to decline the wager, but also why *this* choice is the one that succeeds her torn state. That is, in addition to citing its contribution to her torn state and the associated probability distribution, we could cite it to answer the question, "Why is it specifically *this* choice that ensues, given the probability distribution and everything that accounts for it?" In effect, the Belief would be a *deciding factor*, whose presence gives us something more to say, beyond the probability distribution and everything that accounts for it, about why she ends up choosing as she does, given that distribution. By contrast, we can suppose that the Belief is causally efficacious without supposing that it's a deciding factor or that it gives us something more to say. For if Christine chooses to

decline the wager, the Belief will be among the causes of this choice even if, having fully exhausted its causal influence, it isn't a deciding factor. Thus S2 is stronger than S1.

I submit that the reasonableness of the Second Expectation depends on the possibility of S2 and not merely of S1, for much the same reason that the reasonableness of the First Expectation depends on the reasonableness of the Second Expectation. As I argued in connection with N4, the possibility that Christine's torn choice will comply with the First Expectation is not enough to warrant this expectation; it must also be reasonable to expect it to guide her choice. For apart from such guidance, she has no means by which to secure her compliance, and holding her to this expectation isn't reasonable if she lacks such a means. By itself, the possibility of compliance is too thin a reed on which to hang this normative expectation.

The argument for N5 is directly parallel. Just as the possibility of compliance without the reasonable expectation of guidance isn't enough to warrant the First Expectation, so the possibility of guidance without the possibility of causal influence isn't enough to warrant the Second Expectation. For just as the possibility that Christine will end up complying doesn't amount to a means of bringing it about that she does comply, so the possibility that such guidance will occur doesn't give her a means of bringing it about that such guidance occurs. And without such a means in her torn state, such guidance cannot reasonably be expected of her in that state. Finally, she has no such means in her torn state unless the Belief could causally influence the outcome of that state. Thus the Second Expectation isn't reasonable unless the Belief could causally influence the outcome of her torn state; the mere possibility of its causal efficacy isn't enough.

The argument's final premise is that

- (N6) The Belief could not causally influence which choice actually succeeds Christine's torn state.

Again, by assumption, Christine has the Belief, which is among the states that dispose her to decline the wager. According to N6, there is no possible continuation of her circumstances in which she chooses to decline the wager at least in part because the Belief causally influences the outcome of her torn state. In other words, there is no possible continuation in which the Belief acts as a deciding factor that we can cite to answer the question, "Why is it specifically a choice to decline the wager that actually succeeds Christine's torn state in this (counterfactual) continuation of events?"

That no such continuation is possible is simply a consequence of the assumption that Christine's choice to accept the wager is a truly random outcome of her torn state. Qua truly random, there is nothing more to say, beyond citing the probability distribution and everything that accounts for it, about why Christine makes the specific choice she does as the outcome of her torn state, whichever choice she makes. If the Belief causally influenced the outcome of her torn state, we could cite its causal influence to explain why her torn state has the specific outcome it does, and so there *would* be something more to say: we could say that this specific choice ensues

(at least in part) because the Belief causally influences the outcome. Thus, we cannot consistently suppose both that the Belief causally influences which torn choice ensues and that the ensuing choice is a truly random outcome of Christine's torn state.

If the Belief cannot causally influence the outcome of Christine's torn state, and if we cannot expect her to guide her choice by our normative expectations in her torn state unless such influence is possible, it follows that Christine cannot reasonably be expected in her torn state to guide the ensuing choice by our normative expectations.

But, it might be asked, why must it be possible for the Belief to causally influence the outcome of her actual, torn state? Why construe the guidance requirement in such a way that we hold fixed Christine's actual mental states immediately before she chooses, thereby effectively ensuring that the outcome, qua truly random, won't meet this requirement? Why won't a suitable guidance requirement be met as long as there is *some* point in the run-up to her choice at which the Belief could causally influence the outcome? In that case, although it will be too late for such influence to occur immediately before her actual choice, our libertarian need not accept that such guidance cannot reasonably be expected. For perhaps Christine's earlier susceptibility to this influence is enough to underwrite the Second Expectation, so that this expectation remains reasonable even though the "window" for such influence has closed.

We might put the answer this way. At issue is whether Christine is eligible for moral blame *directly* on the basis of choosing to accept the wager, the choice that actually succeeds her torn state. If she is directly eligible for moral blame on the basis of so choosing, then she is subject to the First Expectation in her torn state, as the maker of the choice that issues from that state. And if she is subject to the First Expectation in her torn state, as the maker of the choice that issues from that state, it must be reasonable to expect her, when she is in that state, to guide her choice by the First Expectation—that is, to choose to decline the wager, as the outcome of this state, because she believes that she ought to decline it. For if we cannot reasonably expect her, while in that state, to guide her choice by the First Expectation, how can we see her as eligible for moral blame directly on the basis of her choice to accept the wager, the choice that actually succeeds that state? Accepting that she cannot reasonably be expected in her torn state to guide her choice by the First Expectation, we cannot see this fact—the fact that a choice to accept the wager succeeds this state—as a direct basis for moral blame; and so we cannot see her as subject to the First Expectation qua maker of this choice. Thus she is not subject to the First Expectation in her torn state unless she can also reasonably be expected, *while in this state*, to guide her choice by it, and this expectation isn't reasonable, in turn, unless the Belief could causally influence the outcome of this state.

From N5 and N6, then, it follows that Christine cannot reasonably be expected to guide her choice by the First Expectation. By N4, she cannot reasonably be held to the First Expectation, and hence, by N3, she is not subject to the First Expectation. And if she is not subject to this expectation concerning her choice, then, by N2, she is not directly morally responsible for choosing to accept the wager, as she actually does. Finally, if she is not directly morally responsible for choosing as she does, then, by N1, she is not directly morally responsible for her

choice simpliciter. Thus, if Christine's choice is a truly random outcome, she isn't directly morally responsible for it. There is no extent to which she is directly morally responsible for so choosing.

In this section, I have proposed an answer to the question of why we aren't morally responsible for choices that are truly random outcomes. The answer doesn't feature the concepts of control, freedom, power, or ability—the concepts we usually find in formulations of the Luck Objection. (For that matter, it doesn't feature the concept *luck* or *chance*.) Instead, it appeals to a guidance requirement for morally responsible agency, a requirement that the relevant choices, qua truly random outcomes, don't meet.

It might be asked why this guidance requirement has gone unnoticed. One possibility is that it's superficially similar to a more familiar requirement for moral responsibility, namely that our choices and actions must be appropriately based on our reasons, so that we can be said to choose and act for those reasons. While the two requirements are superficially similar, the "appropriate causation" requirement is weaker than the guidance requirement. For even if Christine's reasons appropriately cause her choice, whichever choice she makes, she won't satisfy the guidance requirement unless we can reasonably expect her to guide her choice by the First Expectation, and meeting the appropriate causation requirement doesn't make it reasonable to expect such guidance. At the least, going from 'Christine's torn choice is appropriately caused by her reasons' to 'We can reasonably expect Christine to guide her torn choice by the First Expectation' requires a further step. And, I have argued, expecting such guidance of Christine in her torn state isn't reasonable unless the Belief could causally influence the outcome of that state, a requirement that clearly goes beyond the appropriate causation requirement.

Still, it might be asked why the weaker, "appropriate causation" requirement isn't enough. Why can't our libertarian maintain that Christine is morally responsible for randomly choosing to accept the wager, on the grounds that she chooses for her reasons, whichever choice she makes, notwithstanding that we cannot reasonably expect her to guide her choice by the First Expectation?¹⁸ I take up this suggestion in the next section, as part of a libertarian response to the No Guidance Argument.

5. Connecting the Two Stages

In section 2, I argued that supposed exercises of free will would resemble certain nonactions, such as swerving particles and randomly generated numbers, in being truly random outcomes. After considering two libertarian responses in section 3, I set out the No Guidance Argument in the previous section to show that we aren't directly morally responsible for such choices qua truly random outcomes. I submit that pressing these questions in turn—"Are the relevant choices truly random outcomes?" and "Are we morally responsible for them if they are?"—improves on existing formulations of the Luck Objection, in part because the first stage avoids slippery questions about agents' control, powers, and abilities. If libertarians cannot

resist this first stage, they will have to dig in at the second. What are their prospects for doing so?

Recall the first libertarian response from section 2. According to this response, libertarians need not deny that supposed exercises of free will are truly random outcomes as I have defined them, for an exercise of free will just *is* a truly random outcome of the right sort—namely, a mental act of making up one’s mind between two or more genuinely available alternatives. Whichever choice Christine makes, she exercises her ordinary ability to make up her mind, along with her ability to make that particular choice on the basis of her reasons for so choosing.

If Christine’s choice is both a truly random outcome and an exercise of free will, our libertarian might contend that she is directly morally responsible for choosing as she does, despite the No Guidance Argument. In response, I’m prepared to accept the conditional claim that Christine is directly morally responsible for her choice if she exercises free will in choosing as she does. However, I contend that we have much better reason to reject this claim’s consequent than we have to accept its antecedent, given that Christine doesn’t meet the guidance requirement as set out in the previous section.

But suppose that our libertarian disputes the guidance requirement. We touched on this possibility at the end of the previous section. If Christine’s choice is suitably based on her reasons, whichever choice she makes, our libertarian might contend that meeting this “appropriate causation requirement” is enough to sustain the judgment that she is directly morally responsible for choosing as she does, notwithstanding that her choice doesn’t meet the guidance requirement.

Not surprisingly, I find this response highly implausible. After all, our libertarian presumably won’t dispute my claim that Christine is directly morally responsible for her choice only if we can reasonably expect her to choose to decline the wager as the outcome of her torn state. The question then becomes, “Is Christine’s meeting the appropriate causation requirement (whichever choice she makes) enough to support the expectation that she choose to decline the wager, even if we acknowledge that she cannot be expected, in her torn state, to guide her choice by the First Expectation?”

Again, if Christine chooses to decline the wager, she will have guided her choice by this expectation, since the Belief will count as causally efficacious. Our libertarian’s contention, however, is not that she thereby satisfies a plausible guidance requirement; it is, rather, that the occurrence of such guidance is a red herring, since what really matters for moral responsibility is whether her reasons appropriately cause her choice.

In response, I fail to see how meeting the appropriate causation requirement without meeting the guidance requirement could plausibly support our normative expectation concerning Christine’s choice. For unlike the guidance requirement, the appropriate causation requirement doesn’t speak to the question, “When can we reasonably be held to normative expectations concerning our choices and actions?” To see this, suppose that we hold Christine to the First Expectation, morally blaming her for the fact that a choice to accept the wager is the specific outcome of her torn state. If she objects that we’re being unreasonable, given that (by our own admission) she could not have been expected to guide her torn choice by the First

Expectation, and that she had no other means of securing her compliance, it won't do to insist that the expectation remains reasonable because both her actual and her counterfactual choices are appropriately caused by the relevant reasons. For this observation simply doesn't speak to her concern about the reasonableness of holding her to an expectation regarding the outcome of her torn state, when she has no means of securing her compliance with that expectation. To say, "However things turn out, you will have chosen for your reasons," misses the force of the question, "How was I supposed to bring it about in my torn state that the *right* reasons prevailed?" Like the mere possibility of compliance, the possibility of compliance plus appropriate causation is too thin a reed on which to hang our normative expectation regarding Christine's choice.

This is not to deny that our libertarian could in theory propose a different answer to the question, "What is supposed to make it reasonable to hold Christine to the First Expectation in her torn state?" an answer that goes beyond appropriate causation. The point is that the question is a real one, that the above answer falls short, and that it isn't clear what materials our libertarian has for a different answer.

Summing up, our libertarian maintained that exercises of free will are both truly random outcomes and choices that we make for our reasons. If such random choices are exercises of free will, we are presumably morally responsible for them. Against this, I argued that it remains implausible to say that Christine is directly morally responsible for choosing as she does, given the guidance requirement. I then answered the reply that the guidance requirement is superfluous because appropriate causation by the agent's reasons is enough to ground responsibility attributions for torn choices. Thus, libertarians don't seem to have a good response to the No Guidance Argument, which is the second stage of the Two-stage Luck Objection.

As we saw in section 2, event-causal libertarians will have trouble resisting the first stage, which purports to show that supposed exercises of free will are truly random outcomes. For once we have factored all of Christine's causally relevant mental states into the probability distribution, there seems to be nothing more to say about why the one choice or the other actually ensues, given this distribution. By contrast, agent-causal libertarians can distinguish supposed exercises of free will from truly random outcomes precisely on the grounds that the former are agent-caused. But if the Two-stage Luck Objection succeeds against event-causal libertarianism, closely related considerations may well be effective against agent-causal libertarianism. Or so I hope to show in future work.

6. Concluding Remarks

To advance the debate between libertarians and their critics, I have set out a version of the Luck Objection that avoids claims about agents' control, powers, and abilities. The first stage argues only that event-causal libertarians cannot plausibly deny that supposed exercises of free will, though actions, resemble certain nonactions in being truly random outcomes. The second stage denies that we are morally responsible for these actions if they are truly random outcomes. Qua truly random outcomes,

such choices don't meet a plausible guidance requirement for morally responsible agency. Unless they can resist the argument at one of these stages, libertarians who have stopped worrying about the Luck Objection should think again.¹⁹

Notes

¹ Contemporary discussions of this family of objections can be found *inter alia* in van Inwagen 1983, pp. 126–52 and 2000, pp. 14–15; Kane 1996, ch. 7, 1999; Mele 1999, 2006, ch. 5, and 2013; Haji 2001; Clarke 2003, pp. 159–69, 2011; Ekstrom 2003 and 2011; McCall and Lowe 2005; Almeida and Bernstein 2011; Pereboom 2001, chs. 2–3 and 2014, chs. 2–3; Balaguer 2010; Griffith 2010; Shabo 2011, 2013, and 2014; Franklin 2011 and 2012; Levy 2011; Schlosser 2014; and Coffman 2015. It should be noted that not all versions of the objection feature the term 'luck' (or 'chance'), though the family resemblance among different versions is clear.

² While there is no a bright line separating libertarians who see the problem as real but ultimately soluble from those who deny that there is a genuine problem at all, there appears to be a growing skepticism among libertarians (and others) about whether the Luck Objection poses a serious problem for the libertarian enterprise. See for example McCall and Lowe 2005; Griffith 2010; Franklin 2011, 2012, and forthcoming; Ekstrom 2003 and 2011; Balaguer 2010; Almeida and Bernstein 2011; Steward 2012; and Vargas 2012. It should be noted that some of these writers believe that worries about luck are misplaced if (and only if) libertarians adopt an agent-causal framework (O'Connor 2000 and 2011, Griffith 2010, and Steward 2012), while others seem to hold that the worries, at least as currently formulated, are unfounded *tout court* (Ekstrom 2003 and 2011, Balaguer 2010, and Franklin 2011 and 2012).

³ More precisely, the conclusion of the second stage is that we aren't *directly* (or non-derivatively) morally responsible these actions. The restriction to direct responsibility is appropriate since libertarians must say that we are directly morally responsible for at least some such causally undetermined actions.

⁴ In particular, it's often hard to assess the claim that someone has, or lacks, control over a causally undetermined action, since it's often unclear how talk of control is being used. On one construal, to say that someone has the requisite control is simply to affirm that her action meets the relevant (non-epistemic) conditions for moral responsibility, whatever those conditions are. On another possible construal, such talk denotes (or at least gestures toward) a power or ability (e.g. the power to "settle" the outcome) that can be identified independently of—and that can help to ground—responsibility attributions. And there are real questions about how to understand this power or ability. For this reason, defenders of the Luck Objection may do well to avoid couching the challenge to libertarianism in terms of agents' control, power, or abilities.

⁵ Shabo 2011, 2013, and 2014. That said, the version presented here is different in some ways, including ways that help to connect the two stages of the larger argument.

⁶ Since 'anomalous' and 'exceedingly unlikely' are imprecise markers, there will be some fuzziness around the edges of the term 'truly random outcome'. However, this isn't a problem for my purposes, since libertarians must secure free will for actions that don't fall near these edges. As for (b), I doubt that anything would change if we relaxed the condition to include both definite objective probabilities (e.g. 0.46) and probability ranges (e.g. $\sim[0.37-0.46]$), or, alternatively, if we understood the relevant probabilities as irreducibly comparative. For simplicity, however, I have stated (b) in terms of definite probabilities.

⁷ This point will be familiar from discussions of contrastive explanations. Note, however, that contrastive explanations answer questions of the form, "Why did this outcome occur *rather than* that one?" By contrast, the explananda with which I'm concerned have the form, "Of the possible ways things could have turned out, why did they turn out in the specific way they did?" Even if such explanations turn out to be implicitly contrastive in the end, this other formulation is more to the point for my purposes.

⁸ For a different reason, the argument doesn't address noncausal libertarians, either. Nevertheless, if the argument succeeds against event-causal libertarians, I believe that it can be extended to event- and agent-causal libertarians as well, as I hope to show in future work.

⁹Except, of course, in the trivial sense that we don't meet the "control" or ("freedom-relevant") condition for moral responsibility, whatever this condition is presumed to be.

¹⁰See especially Balaguer 2010 and Franklin 2012. A noteworthy feature of this reductive approach is that it doesn't require any resources that aren't employed by compatibilist accounts. In this respect, Balaguer's and Franklin's responses to the Luck Objection are similar to Mele's Daring Soft Libertarian response (Mele 2006, ch. 5). See also Fischer 2012, ch. 6. Thanks to an anonymous referee for prompting me to note the parallel with Mele.

¹¹Kane's view is illustrative here. Kane wants to say that we have "control over" which choice we make in exercises of free will (see for example Kane 1999, pp. 236–39, and forthcoming), and he would presumably deny that we have such control over truly random outcomes as I've defined them (see Kane 2014, pp. 55–57). Thus, unlike libertarians who favor the first response, he must hold that there is a crucial difference between free will and choices like Christine's. Yet Kane hasn't explicitly said whether "control over" the outcome requires a further fact (one that is absent in an otherwise similar random-outcome scenario).

¹²Perhaps it will also include some agent-causal libertarians who wonder whether *the agent-causing of a choice* could be a truly random outcome, even if an agent-caused choice could not be.

¹³I take no stand on whether this mental state (the state of believing that she should decline the wager) could by itself causally influence which choice she makes, or whether it could do so only in conjunction with a distinct motivational state.

¹⁴Without the restriction to self-governing agents, N4 would arguably be susceptible to preemption-type counterexamples, similar to Frankfurt's (1969) famous counterexample to the Principle of Alternative Possibilities. I am indebted to an anonymous referee for comments that prompted me to spell out this part of my argument more fully.

¹⁵To be clear, the scope of the relevant belief is that *the agent ought to take a particular course*, viz. the course that is specified by the expectation to which we'd be holding her. The scope of the belief is not: *that the agent ought to take the course that we expect her to take*. (Our expectation helps to identify the content of the belief, without being part of that content.)

¹⁶To anticipate, I consider how a libertarian might seek to resist the guidance requirement in section 4.

¹⁷It is here that the affinity between the Two-stage Luck Objection and the Disappearing Agent Objection to event-causal libertarianism, as formulated by Pereboom (2014, chs. 2 and 3), is clearest. I compare the two arguments in work in progress, where I hope to show that the No Guidance Argument should get more traction with minimalist libertarians like Balaguer and Franklin.

¹⁸Thanks are owed to an anonymous referee for prompting me to address this response.

¹⁹I want to thank two anonymous referees and an editor at *Noûs* for valuable comments that led to numerous improvements in the manuscript. For helpful conversation about the No Guidance Argument, I'm grateful to Gunnar Björnsson, Michael McKenna, Derk Pereboom, and Jim Stone.

References

- Almeida, M. and M. Bernstein. (2011). "Rollbacks, Endorsements, and Indeterminism," in R. Kane (ed.), *The Oxford Handbook of Free Will*, 2nd edn. (New York: Oxford University Press), pp. 484–95.
- Balaguer, M. (2010). *Free Will as an Open Scientific Problem* (Cambridge: MIT Press).
- Buchak, L. (2013). "Free Acts and Chance: Why the Rollback Argument Fails," *Philosophical Quarterly* 63, pp. 20–28.
- Clarke, R. (2003). *Libertarian Accounts of Free Will* (New York: Oxford University Press).
- Clarke, R. (2011). "Alternatives for Libertarians," in R. Kane (ed.), *The Oxford Handbook of Free Will*, 2nd edn. (New York: Oxford University Press), pp. 329–48.
- Coffman, E. J. (2015). *Luck: Its Nature and Significance for Human Knowledge and Agency* (New York: Palgrave Macmillan).
- Ekstrom, L. (2003). "Free Will, Chance, and Mystery," *Philosophical Studies* 113, pp. 153–80.
- Ekstrom, L. (2011). "Free Will is Not a Mystery," in R. Kane (ed.), *The Oxford Handbook of Free Will*, 2nd edn. (New York: Oxford University Press), pp. 366–80.

- Fischer, J. M. (2012). *Deep Control: Essays on Free Will and Values* (New York: Oxford University Press).
- Fischer, J. M. (2014). "Toward a Solution to the Luck Problem," in D. Palmer (ed.), *Libertarian Free Will* (New York: Oxford University Press), pp. 52–68.
- Frankfurt, H. (1969). "Alternate Possibilities and Moral Responsibility," *Journal of Philosophy* 66, pp. 829–39.
- Franklin, C. E. (2011). "Farewell to the Luck (and *Mind*) Argument," *Philosophical Studies* 156, pp. 199–230.
- Franklin, C. E. (2012). "The Assimilation Argument and the Rollback Argument," *Pacific Philosophical Quarterly* 93, pp. 395–416.
- Franklin, C. E. (forthcoming). *A Minimal Libertarianism: Free Will and the Promise of Reduction* (New York: Oxford University Press).
- Griffith, M. (2010). "Why Agent-caused Actions Are Not Lucky," *American Philosophical Quarterly* 47, pp. 43–56.
- Haji, I. (2001). "Control Conundrums: Modest Libertarianism, Responsibility, and Explanation," *Pacific Philosophical Quarterly* 82, pp. 178–200.
- Kane, R. (1996). *The Significance of Free Will* (New York: Oxford University Press).
- Kane, R. (1999). "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism," *Journal of Philosophy* 96, pp. 217–40.
- Kane, R. (2014). "Torn Decisions, Luck, and Libertarian Free Will: Comments on Balaguer's Free Will as an Open Scientific Problem," *Philosophical Studies* 169, pp. 51–58.
- Kane, R. (forthcoming). "The Complex Tapestry of Free Will: Striving, Indeterminism, and Volitional Streams," *Synthese*.
- Levy, N. (2011). *Hard Luck* (Oxford: Oxford University Press).
- McCall, S. and E. J. Lowe. (2005). "Indeterminist Free Will," *Philosophy and Phenomenological Research* 70, pp. 681–90.
- Mele, A. (1999). "Ultimate Responsibility and Dumb Luck," *Social Philosophy and Policy* 16, pp. 274–93.
- Mele, A. (2006). *Free Will and Luck* (New York: Oxford University Press).
- Mele, A. (2013). "Moral Responsibility and the Continuation Problem," *Philosophical Studies* 162, pp. 237–55.
- O'Connor, T. (2000). *Persons and Causes: The Metaphysics of Free Will* (New York: Oxford University Press).
- O'Connor, T. (2011). "Agent-causal Theories of Freedom," in R. Kane (ed.), *The Oxford Handbook of Free Will*, 2nd edn. (New York: Oxford University Press), pp. 309–28.
- Pereboom, D. (2001). *Living Without Free Will* (Cambridge: Cambridge University Press).
- Pereboom, D. (2014). *Free Will, Agency, and Meaning in Life* (New York: Oxford University Press).
- Schlosser, M. (2014). "The Luck Argument Against Event-causal Libertarianism: It is Here to Stay," *Philosophical Studies* 167, pp. 375–85.
- Shabo, S. (2011). "Why Free Will Remains a Mystery," *Pacific Philosophical Quarterly* 92, pp. 105–25.
- Shabo, S. (2013). "Free Will and Mystery: Looking Past the *Mind* Argument," *Philosophical Studies* 162, pp. 291–307.
- Shabo, S. (2014). "Assimilations and Rollbacks: Two Arguments Against Libertarianism Defended," *Philosophia* 42, pp. 151–72.
- Steward, H. (2012). *A Metaphysics for Freedom* (Oxford: Oxford University Press).
- van Inwagen, P. (1983). *An Essay on Free Will* (Oxford: Clarendon).
- van Inwagen, P. (2000). "Why Free Will is a Mystery," *Philosophical Perspectives* 14, pp. 1–20.
- Vargas, M. (2012). "Why the Problem of Free Will Isn't," *Philosophical Issues* 22, pp. 419–36.