

Libertarianism and Skepticism about Free Will: Some Arguments against Both

Manuel Vargas

University of San Francisco

On one way of putting things, incompatibilism is the view that in some important sense free will (and/or moral responsibility) is incompatible with determinism. Incompatibilism is typically taken to come in two species: libertarianism, which holds that we are free and responsible (and correspondingly, that determinism does not hold), and skeptical incompatibilism.¹ The latter includes views such as hard determinism, which hold that we are not free (and/or responsible) and views that argue that free will is incompatible with both determinism and indeterminism, among others. In this paper, I attempt to provide positive arguments against *both* of the primary strands of incompatibilism.

The first aim of this paper is to take some steps toward filling in an argument that is often mentioned but seldom developed in any detail—the argument that libertarianism is a scientifically implausible view. I say “take some steps” because I think the considerations I muster (at most) favor a less ambitious relative of that argument. The less ambitious claim I hope to motivate is that there is little reason to believe that extant libertarian accounts satisfy a standard of naturalistic *plausibility*, even if they do satisfy a standard of naturalistic *compatibility*.

The second aim of this paper is to argue against skepticism about free will without denying the presence of incompatibilist intuitions. Indeed, I am inclined to think that many of us *do* have incompatibilist intuitions and that they reflect an important aspect of our self-conception. What I endeavor to provide are considerations for thinking that neither the shortcomings of libertarianism nor the difficulties of standard arguments for free will skepticism are sufficient for embracing skepticism about free will and/or moral responsibility.

I start with some methodological considerations about the aim of theorizing about free will. I then argue for the comparative implausibility of libertarianism, followed by an argument against free will skepticism. The last section of the paper considers the alternatives that remain for those who feel an impulse toward incompatibilism but accept skepticism about libertarianism and skepticism about skepticism about free will.

I. METHODOLOGICAL CONSIDERATIONS

On at least one conception, the philosophical endeavor is something like the search for probable truths in domains in which we lack any reliable methodology for determining what the truth is. On this conception of things, it is natural to suppose that the closer we get to attaining the aim of the inquiry, the more demanding the standards become, perhaps even involving some degree of predictive success or integration into our epistemically best explanatory frameworks. And, in those rare cases where we are successful enough, the endeavor may even cease to count as philosophy in any straightforward sense.² During the early stages of philosophical inquiry, though, we expect comparatively little. Indeed, it will often be enough to show that some controversial view is possible. And, it will be a sign of progress when we come to understand how a view we once thought unthinkable or impossible to intelligently defend in widely acceptable terms is intelligible and possible, after all.

Considered in this light, libertarianism about free will (the view that we have free will but that it is incompatible with determinism) has made significant progress in the past few decades. We now have some sense about how it could turn out to be true, even if many of us do not think it is true or is likely to be true. And, I think the consensus is that this is an important difference between current debates and the debates that were had even little more than twenty years ago. The sense one gets from examining anthologies of that time is that no one had the faintest idea what a non-ad hoc or even nonincredible picture of libertarianism would amount to—even when

there were compelling arguments for thinking that one had to be true. For most of the twentieth century, the stench of souls, Cartesian dualism, Kantian noumena, and other forms of what P. F. Strawson called “panicky metaphysics”³ seemed to cling to the very idea of libertarianism. In the dialectical economy of that earlier period, the specter of metaphysics in the pejorative sense appears to have provided much of the impetus behind compatibilism. Even today, there is still some pull to the Familiar Argument: Since we are *obviously* responsible, if incompatibilism is committed to that sort of [insert your favorite implausible metaphysics], compatibilism *must* be true!⁴ Since there was no worked-out account of how libertarianism’s ontological commitments could be compatible with the emerging scientific understanding of the world, about the only thing that could be said in favor of libertarianism by its defenders was the controversial claim that it was the only way to capture common-sense thinking about free will and moral responsibility.⁵

Though it is somewhat anachronistic, the point can be usefully cast in the following fashion: until relatively recently, almost no one had any sense of how libertarianism could, in some non-ad-hoc way, satisfy a *standard of naturalistic compatibility*. By ‘naturalistic compatibility’ I mean nothing more than compatibility with an independently acquired, broadly scientific conception of the universe, especially the parts we inhabit.⁶ There are, of course, philosophers who reject naturalism in one or another form.⁷ However, acceptance of the standard I am concerned with—compatibility with a broadly scientific worldview—has widespread acceptance, even if many forms of naturalism do not. The motivation for it is relatively straightforward: a philosophical theory that is in tension with our best science is a philosophical theory that is in tension with our best pieces of knowledge.

For better or for worse, virtually all philosophically serious positive accounts of free will accept a standard of naturalistic compatibility.⁸ In our current philosophical climate, a theory that does not is simply a nonstarter—unless it is introduced to be rejected for one reason or another.⁹ Fortunately, at least for those with libertarian inclinations, we now have a range of accounts that

seem to satisfy the standard of naturalistic compatibility.¹⁰ From the more metaphysically modest forms of libertarianism, such as those offered by Robert Kane,¹¹ Laura Ekstrom,¹² and Al Mele,¹³ to more metaphysically adventurous, yet sophisticated, versions of agent-causation (such as those developed by Randolph Clarke¹⁴ and Tim O'Connor¹⁵), contemporary libertarian theories have shown how libertarianism need not (and perhaps should not) be committed to the troublesome metaphysics of older versions of libertarianism.

Libertarianism's progress is clear: contemporary libertarian accounts show the conceivability of something that once seemed inconceivable under a broadly scientific understanding of the world. Contemporary libertarianism cannot be dismissed as just another mystical piece of hokum by tender-minded philosophers. Congratulations are in order to libertarians for having made genuine progress on a hard topic, one where some seem to have thought no progress was even possible.¹⁶

New achievements often bring with them new challenges, and it is no different in the case of libertarianism. Although contemporary libertarian theories may satisfy a standard of naturalistic compatibility, they do not satisfy a more demanding standard of naturalistic *plausibility*. On a standard of naturalistic plausibility, it is not enough that the theory simply avoid contradiction with a scientific worldview. On this standard, we demand a theory that is both compatible with a scientific worldview *and* that the balance of known truth-relevant considerations would be sufficient to lead a group of informed, well-reasoning, and disinterested persons to think the theory is plausible.

The language of plausibility may invite some confusions I wish to forestall. I do not mean plausibility in the sense whereby something is plausible only if we think it is likely to be true. Nor do I mean plausible in the sense of something being more likely to be true than any of the other known alternatives, even if no particular view (considered by itself) strikes us as likely to be true. The sense of plausibility I am interested in is one where the balance of known truth-

relevant considerations could, via an appropriate deliberative path, lead a group of informed, well-reasoning, and disinterested persons to believe the theory.

There is good reason to care about satisfying this more demanding standard. Given the picture of philosophical labors I began with—generating probable truths in a domain in which we lack reliable methods for determining the truth—this raising of the dialectical bar is something of an inevitable outcome of libertarianism's recent success. Since we are in the business of generating probable truths (that is, claims that we have good reason to think are likely true), it will not be enough to show that some theory is not impossible, given our best knowledge. This is, after all, the position contemporary libertarians are in right now. Rather, we should endeavor to push our theories ever closer to being recognizably probable truths. The shift from naturalistic compatibility to naturalistic plausibility is an important step in that direction.

One might wonder whether acceptance of this standard will make any difference at all in our evaluation of existing libertarians. I think that it does make a difference. Few libertarians have offered any reason for us to suppose that their account must be true, apart from our wishing it were so, and what data there is about how the mind works does little to suggest that any account of libertarianism is true. This is the subject of section 2 of this paper.

Alternately, one might wonder whether we should raise the bar even further. Why not go ahead and demand more than plausibility? Why not demand the thing we are actually hoping to get—a theory we have reason to think is, in fact, probably true?

I am inclined to think we ought not raise the bar that high. First, we should be clear about what a standard of probable truth requires. For something to be a probable truth in the relevant sense, it needs to possess something more than a favorable estimation of subjective probability. Whether it needs to be objectively probably true or intersubjectively probably true and how to cash out these and relevant other notions is beyond the scope of this paper. What is not enough, though, is for a particular theorist to believe of his or her own theory that it is probably true. Many, perhaps most theories can satisfy this standard. Second, there is a sense in which, at least

for many of us, the bar is always already raised to the standard of naturalistic probable truth. Inasmuch as we are committed to the picture of philosophy described at the start of this section, the ultimate appropriate aspiration for any philosophical theorizing is that it eventually contributes to (or even constitutes) a theory that reliably gets at truths in some domain. What is at stake, then, is not some master end for philosophical theorizing. Rather, what is at stake is how we are to conduct the more immediate-term evaluation of various philosophical proposals. This latter project clearly benefits from an incrementalist approach. Raising the standard too high will discourage explorations that would bear some fruit if given the opportunity. Conversely, allowing theories to rest on their proverbial laurels when they have achieved success at some relatively low-level standard of evaluation is not conducive to achieving our aims, either. Hence, when theories become successful at (for example) meeting a standard of naturalistic compatibility, it is time to raise the bar. And, when theories become successful at meeting a standard of naturalistic plausibility, it will be time to raise the bar again.

A different kind of concern about the proposal I am offering is this: one might worry that we simply are not in a position to evaluate the naturalistic plausibility of any given theory in this domain. For example, some have argued that since we do not know how future science will work out, and for all we know, it may favor libertarianism, libertarianism is no more or less scientifically plausible than compatibilism.¹⁷

We should reject this line of reasoning for two reasons. First, while it is true that we do not know how future science will work out, we can know what current science says about various issues relevant to the theory of agency. And, as I will go on to argue, what evidence there is does not seem to favor our best philosophical accounts of libertarianism. Second, even in cases where scientific research have nothing to say, we can at least approximately measure the demands we are putting on future theories. All things being equal, a theory with fewer demands on the outcomes of future science ought to be treated as more plausible than a theory with greater demands on the outcomes of future science, simply because there are more ways for a

demanding theory to turn out to be false. This is not to say that we want theories that cannot be falsified, or that we want theories that make minimal commitments. On the contrary, the scramble toward probable truths is impossible without increasingly refined and ontologically specific theories. Rather, the point is that when we are evaluating competing naturalistically compatible theories with roughly equal virtues, we should regard with greater skepticism the theory with more extravagant commitments. Given the somewhat unspectacular history of metaphysical speculation about human agency, we must think carefully about the extent to which our theories are plausible or not in light of what we can and do have reason to expect from current and future science.

There is nearly always an unremarked upon elephant that lurks in rooms where philosophers discuss free will. In this instance, the elephant may be more difficult to ignore. The elephant is the role of religion in motivating and sustaining various libertarian accounts. It would, I think, be revealing to do a survey of the religious beliefs of contemporary libertarians and compatibilists. My guess is that we would learn that a disproportionate number—perhaps even most—libertarians are religious and, especially, Christian. I suspect we would also learn that the overwhelming majority of compatibilists are atheist or agnostic. (I do not have a guess about the religious predilections of the various stripes of skeptical incompatibilists.) Even so, it may not be obvious why the religious beliefs of particular philosophers should matter. We could think that the arguments of the various partisans in free will debates should be judged on their merits, irrespective of whether or not they have religious motivation. However, systematically ignoring the role of religion can remove important considerations from view, considerations that affect how we think about, argue, and evaluate the various philosophical possibilities. Elsewhere, I have argued that one key to understanding the intractability of free will debates is to recognize the presence of methodological differences between those with primarily metaphysical concerns and those with primarily normative concerns.¹⁸ I now also think that understanding the difference religion can make may be a key to understanding some important methodological differences

between religious libertarians and their interlocutors. Though one might be a libertarian who is religious (or, for that matter, a compatibilist who is religious), a *religious libertarian* in my sense is one who, antecedent to and perhaps independent of philosophical inquiry, is committed to a strong belief in a particular divine moral order that requires a strong notion of human freedom. In the doxastic economy of the religious libertarian, the libertarianism is inextricably tied to a religious framework. Consequently, views such as hard incompatibilism and compatibilism will be viewed in light of those commitments, and to the extent that those commitments are strong, hard incompatibilism and compatibilism will be viewed as unpalatable options. The task for religious libertarians is to explain *how* their preferred order could be. *Whether* we have libertarian free will is typically taken to be settled prior to and independent of the philosophical discussions that make up the contemporary scholarly literature on free will. Hence, if we want to understand the dialectic of certain aspects of the free will debate, we cannot do so independent of these considerations.¹⁹

If I am right about this, none of what follows has much chance of persuading religious libertarians to give up their libertarianism. Fortunately, the argument I offer here is not intended to accomplish that much. Rather, the aim of the argument is to convince anyone who already accepts a standard of naturalistic compatibility to accept an even more demanding standard of naturalistic plausibility. What special pressure this puts on religious libertarians (for example, whether acceptance of this standard and a subsequent inability to meet it would spell special trouble for the possibility of shared discourse between religious libertarians and those who are not) merits more reflection than will be offered in this paper.

II. FROM COMPATIBILITY TO PLAUSIBILITY

A standard of naturalistic compatibility is too permissive. It does not give us any way to distinguish clearly inferior theories from better and more plausible theories among those that are

not flatly in opposition to a broadly scientific view of the world. A standard of naturalistic plausibility, however, gives us a useful and principled metric for doing just that. Recall that the standard of naturalistic plausibility is one that holds that a theory is naturalistically plausible if it is both compatible with a broadly scientific picture of the world *and* that the balance of known truth-relevant considerations would be sufficient to lead a group of informed, well-reasoning, and disinterested persons to accept the theory. To the extent that an account has unsupported theoretical demands, or commitments for which we lack independent evidential support, we can expect that disinterested persons will rightly think it implausible.

The upshot is that if we are to accept a theory of freedom and responsibility, it had better meet the standard of naturalistic plausibility. However, this raises a serious difficulty for virtually all libertarian theories. In every instance of libertarianism, whether uncaused-event libertarianism, agent-causal libertarianism, or event-causal libertarianism, we are asked to accept that the world is constructed in some fashion for which we have no other reason for thinking it is built that way other than it would be (if they are right) a felicitous alignment between our (likely) culturally and historically contingent common-sense metaphysics and the way of the world. In other words, we are being asked to take on commitments for which we lack independent evidential support.

The point is a comparative one. Given that virtually all philosophy is somewhat speculative, we will almost inevitably be invited to take on commitments that outstrip our evidential basis. The point is only that we ought not multiply commitments beyond necessity, and the more those commitments are demands, the more unnecessary they become. Faced with a choice between an account that preserves some feature of common sense at the cost of implausibility, and an account that gives up some feature of common sense without implausibility, it seems to me that philosophers committed to the pursuit of truth ought to favor the latter and discount the former.

In response, some libertarians might be tempted to emphasize the role of considerations other than those that play a specifically truth-supporting role (e.g., the dignity or value of libertarian free will). However, this sort of move seems spurious given the assumed conception of philosophy as an endeavor that aims at probably true theories. Unless we are given reason to think that non-truth-supporting considerations are relevant to the probable truths of a theory, it seems altogether irrelevant to raise such considerations. However, even supposing that these non-truth-supporting considerations are somehow relevant or important for how we evaluate a theory, the difficulty for the libertarian will be to explain why these same considerations cannot be met by a compatibilist construal of the salient conditions.

A different sort of reply that may appeal to some libertarians is to argue that evidence from introspection favors libertarianism, and that as long as libertarianism is not ruled out by a scientific worldview, the introspective evidence is sufficient to ground confidence in libertarianism. This line of response is even less promising, to my mind. That we understand ourselves as having a significant sort of freedom is not evidence that we do have that freedom, or even that the freedom we presuppose is really libertarian.²⁰ This is not the place to canvass the literature against introspective evidence for our freedom, but it should suffice to note that anyone familiar with contemporary social psychology will regard introspective evidence, especially with respect to the explanations we offer for our own mental processes, with considerable skepticism.²¹ This is not to say the introspection could not, in principle, provide us with some evidence of one or another sort of freedom. Rather, it is only to say that the burden is on those who think introspection is a reliable guide to the metaphysics of agency to explain why we should suppose that introspection about free will will tell us anything about the metaphysics of agency. To my mind, libertarians have offered no such argument.

So, in contrast to compatibilist accounts, libertarianism will always have at least one additional theoretical commitment—it must postulate the presence of indeterminism (one way or another) in mental processes, and this commitment outstrips any evidence we have for its being

true. As acute as this problem may be, it can become even worse when libertarian theories are sophisticated enough to make clear what sorts of commitments they intend to take on. This is because one thing we can learn is that a theory's commitments run *counter* to the going views in some or another field of science.

Consider Robert Kane's admirably developed account of libertarianism. One of the many virtues of his account is that he is very explicit about how he imagines things would need to turn out for his account to be correct. He holds that free will is located in "self-forming actions" (SFAs), which are actions or willings that are indeterministic, intentional, voluntary, and endorsed by the agent in circumstances where he or she faces a choice between competing motivations. The basic picture is this:

There is a tension and uncertainty in our minds at such times of inner conflict that is reflected in appropriate regions of our brains by movement away from thermodynamic equilibrium—in short, a kind of stirring up of chaos in the brain that makes it sensitive to micro-indeterminacies at the neuronal level. As a result, the uncertainty and inner tension we feel at such soul-searching moments of self-formation is reflected in the indeterminacy of our neural processes themselves. What is experienced phenomenologically as uncertainty corresponds physically to the opening of a window of opportunity that temporarily screens off complete determination by the past. . . .

. . . When we decide in such circumstances, and the indeterminate efforts we are making become determinate choices, we *make* one set of competing reasons or motives prevail over the others then and there *by deciding*.²²

Kane supposes that there are at least two recurrent, connected, and competing "neural networks" in an agent, with each network being characterizable as containing as its input a desire, motivation, or consideration and as its output some choice that satisfies the desire. The networks are competing because the agent cannot choose to satisfy both desires, and the satisfaction of one precludes the satisfaction of the other. Since the networks are connected and chaotic, according to Kane, the conflict between them makes them susceptible to lower-level indeterminacies. When the agent decides on one option over the other, this corresponds to one of

the neural pathways reaching an activation threshold, overcoming the indeterminism of the other.²³

The language of chaos ought not mislead anyone—chaotic systems are strictly deterministic, albeit systems that are virtually unpredictable. However, Kane is not supposing that the chaos is indeterministic. Rather, the idea seems to be that under particular conditions there are indeterminacies at the level of neurons (i.e., cells) that contribute to which network of the neurons will reach “activation threshold,” thus settling what the agent will do. The main difficulty of this account, and indeed any account that postulates indeterminacies in the brain, is that there simply are no accepted scientific models of indeterministic events in the brain. What models we do have tend to be deterministic, and what evidence there is concerning indeterminacies in the brain weighs against indeterministic pictures of the mind. This is true whether we imagine the brain events to be irreducibly indeterministic, or whether we suppose that they amplify or reflect lower-level indeterminacies.

One of the few attempts to show how indeterminacies might find purchase in the brain is Penrose and Hameroff’s highly speculative account of quantum-sensitive microtubules. But philosophers, mathematicians, and neuroscientists have vigorously criticized this account on a number of grounds.²⁴ Although the state of contemporary brain science does not rule out the possibility of indeterministic events in the brain, what we do know seems to weigh against it. As the neuroscientist and philosopher Henrik Walter has concluded, “to date there is no solid empirical evidence that local quantum phenomena play a role in neurons, and that there are good arguments to the contrary.”²⁵ What a Kane-style account saddles us with, then, is substantial demands on how future science must turn out, but also demands that the state of the relevant science seems to weigh against.

Some might be inclined to conclude that this is so much the worse for neuroscience. Although neuroscience is in its relative infancy as far as sciences go, such dismissiveness requires an overly optimistic assessment of the epistemic credentials of philosophy and

conceptual analysis for telling us about the construction of the brain, in comparison to the more empirically oriented tools of neuroscientists. However, one could perhaps more justifiably complain that neuroscientists are working with unimaginative, under-informed, or overly simplistic models of how the brain operates. And, so the complaint might go, if they were working with the right models, then they might interpret the available data and evidence in a different fashion, one that might render plausible indeterministic models of the brain.

This may well be right. The vast majority of contemporary neuroscientists are working with highly impoverished conceptions of agency, and they certainly would do well to attend to some of the distinctions and concepts developed by philosophers who work on aspects of agency. Notably, however, what they are insensitive to is not libertarianism, but the possibility of compatibilism.²⁶ But even if we suppose that the field could be reshaped in ways that are friendlier to current libertarian views, we have no reason to suppose, here and now, that such a change will be forthcoming. In turn, this means that philosophers have little justification for accepting models of libertarianism that rely on the existence of indeterminacies that show up in the brain. The lesson to be learned by thinking about Kane's account is that until neuroscientists propose workable indeterministic models of the brain, or until libertarians undertake serious neuroscientific work, we should remain skeptical about purely speculative claims as to how the neuroscience must work out.

We can sum things up in the following fashion: libertarians are asking us to take on commitments that outstrip our evidential basis in defense of various purported features of common sense. This, by itself, is comparatively troubling given the existence of compatibilist models of agency that make fewer demands on how the world must turn out to be. However, the particular commitments of libertarianism seem especially troubling to the extent to which they require indeterminism to show up in very particular places, and at least on some models, in very particular places in the brain. Given that there are no credible scientific model of indeterminacies in the brain, and given that there are reasons internal to neuroscience for thinking that the brain is

not so organized, libertarianism of especially this sort—though arguably, of any sort that holds that indeterministic mental processes supervene on the brain—will fail to meet a standard of naturalistic plausibility.

III. CONSEQUENCES AND ALTERNATIVES

A satisfactory theory of free and responsible agency must do more than satisfy a standard of naturalistic compatibility—it must satisfy a standard of naturalistic plausibility. However, there is little to suggest that libertarian theories currently satisfy this standard. In contrast, standard forms of compatibilism seem somewhat less likely to run afoul of the standard. This is not to say that compatibilist theories are without their own difficulties, or that they do not make demands of their own. Rather, I have endeavored to argue that libertarianism is saddled with additional demands, and that those additional demands render those theories comparatively less plausible than compatibilist theories, at minimum, and at maximum they render those theories flatly implausible given the state of contemporary science. Since there is good reason to accept a standard of naturalistic plausibility,²⁷ and since libertarian theories appear to fall considerably short of this standard, it seems that we should regard libertarian theories as implausible.²⁸

Suppose that we accept all of the preceding points. Some incompatibilists will be inclined to keep up the good fight, attempting to show that a standard of naturalistic plausibility, or something very much like it, can be satisfied, despite libertarianism's comparatively poor track record of naturalistic plausibility over the past two millennia. Some might embrace "mysterianism," the view that we have free will, but that it is a mystery how we have it. This is an unattractive option. One consequence of mysterianism seems to be that the justification of our differential treatment of others (especially blame and punishment) depends on a mysterious notion. This leaves us in the unfortunate position of insisting that people are free and responsible without providing any accessible or intelligible basis on which to justify our differential

treatment of others. This is hardly the stuff that provides us with principled guidance and the tools for collective deliberation in a pluralistic society.

What of free will skepticism? If one were moved by the various general arguments for incompatibilism, but agreed that libertarianism is comparatively implausible, would some form of skepticism about free will be the appropriate position to embrace? In the next section I argue the answer is no.

IV. SKEPTICISM ABOUT FREE WILL

Though the details differ, skeptics about free will typically offer arguments with the following basic structure: First, we are given a characterization of what it is to have free will (or some capacity or form of agency required for moral responsibility) that is supposed to be consistent with our untutored intuitions about it.²⁹ Then, we are given an argument for why it is implausible or even impossible to have some feature of that characterization. Finally, we are then urged to accept the skeptical conclusion that we do not have free will or moral responsibility.

Galen Strawson's "Basic Argument" is representative of the genre, though nearly any argument for free will skepticism will do. Here is one version of the argument he offers:

1. You do what you do, in any situation in which you find yourself, because of the way you are. So,
2. To be truly morally responsible for what you do you must be truly responsible for the way you are—at least in certain crucial mental respects. But,
3. You cannot be truly responsible for the way you are (because you would have had to intentionally brought it about that you are the way you are, where the basis of your bringing it about was itself intentionally brought about, on a basis that you intentionally brought about, on a basis that . . .), so you cannot be truly responsible for what you do.³⁰

All the elements of standard arguments for skepticism about free will are present: we are given a characterization of some requirement that an agent is supposed to satisfy to be morally

responsible, we are then given an argument for why that characterization cannot be realized, and then we are urged to accept the skeptical conclusion.³¹

To see what is wrong with arguments of this form, we can begin by giving the skeptic the premises. Let us suppose that the skeptic's characterization of the requirements for free will or moral responsibility strike us as plausible. Let us also suppose the skeptic has successfully shown that there is good reason to think that this characterization is impossible to have or that it is implausible to think that we have it. Even so, the conclusion does not follow. All that follows is this: if free will is like we imagine it to be, then we do not have it. However, we are not entitled to conclude from this that we lack free will, for free will might be somewhat different from what we imagine it to be. What the skeptic requires, then, is a further premise or two—that free will is just as we imagine it to be, and that it cannot be otherwise. Without these elements, the skeptical conclusion is stymied by the possibility that free will is or could come to be somewhat different than the skeptic's characterization of it.

Some will protest that the additional premises are obvious or otherwise reasonable to grant the skeptic. But given what is usually cited as being at stake—our self-image, practices of praise and blame, and the deservingness of different ways of treating one another—we should be cautious about surrendering these premises too easily. All too often we have found that the nature of the world does not neatly reflect our conception of it, and there does not seem to be any special reason to suppose that free will in the world—whatever that amounts to—will be an instance of a fortuitous alignment between conception and reality. On the contrary, there is reason to suppose that free will, if it does exist, will *not* be precisely as we imagined it. Given the particular sociocultural history of the concept, and in particular, the role it played in Christian theology and pre-scientific conceptions of the self, it seems unduly optimistic to suppose that this particular culturally inherited concept will have come down to us in a form that is smoothly compatible with a contemporary scientific view of the world. Indeed, nearly *any* significant concept—physical, moral, or otherwise—that has a long enough history is unlikely to survive

unrevised in the face of growing knowledge about the world. Given that the notion of simultaneity proper to physics, our moral notions of what constitute virtues, and our conception of marriage³² all have been subject to revisions in various ways, we need some special reason to suppose that free will is different from these cases. Since these skeptics have not (thus far) offered such a reason, there is no reason for us to suppose that the troublesome feature of free will identified in a given skeptical argument is really a(n immutable) part of free will proper, as opposed to a contingent feature of how we currently, and perhaps erroneously, think about free will.

This point holds irrespective of whether or not one favors an internalist conception of the semantics of free will, whereby the meaning of ‘free will’ is decided solely by our linguistic or conceptual practices, or whether one favors an externalist conception whereby the meaning of the term is in part fixed by some contribution of the world, apart from our linguistic or conceptual practices. For the externalist, the issue is simple: skeptics have given us no reason to suppose that the thing itself must have the skepticism-inducing characteristic of our concept. For the internalist, the issue is whether there is any reason to think that the characteristic of the concept identified by the skeptic is something essential and unrevisable as such that anything that lacks it could not, as a matter of principle, count as free will.

The conclusion that the identified aspect of free will is essential and unrevisable cannot be earned by mere assertion. Without some principled reason to suppose that free will cannot be other than as we conceive of it, or some explanation of why free will is different from so many other concepts that have undergone change as we learned more about the world, it would be sheer prejudice to insist on so tight a connection between our current historical conception of free will and the world. *Prima facie*, any argument for the essential indispensability of some aspect of our concept of free will would seem to be an extraordinarily difficult thing to show in light of what we do know about our flexibility in making sometimes drastic changes in our concepts.

At least in conversation, some philosophers are tempted to respond with a philosophical throwing up of the hands: Since we do not have any other way to get at the nature of free will other than conceptual analysis,³³ free will must be as conceptual analysis reveals it to be, these philosophers say. What makes this a bad argument—outside of something like a Kantian project that attempts to describe the necessary character of the phenomena, perhaps—is that it presupposes that the nature of the world is to be read off of our current epistemic limitations. That we do not currently know how to “get at” free will in some fashion not completely dependent on conceptual analysis does not mean that we might not someday have the means to do so. Of course, we may never have the means to do so. And, in fact, perhaps there is no way to ever, even in principle, understand the nature of free will apart from conceptual or semantic analysis. But, this does not itself mean that the nature of free will is settled once we have done conceptual analysis. Conceptual analysis may guide our attempts to understand free will by giving us an account of what to look for in the world, but we ought not suppose that the world cannot teach us things that lead us to revise our conception of free will, any more than we should suppose that discoveries in social psychology cannot change or transform our understanding of ourselves.³⁴

In short, the problem the skeptic faces is this: the skeptical argument has a gap in it that requires some further premises, premises which are not themselves obviously true. Given what is at stake, and given that the premises are not obviously true, it would be a mistake to accept the skeptical conclusion. Therefore, the challenge for the skeptic is to close the gap, either by showing that the further premises are true, or that there is some other route to the skeptical conclusion. Since I am unaware of any *argument* in favor of the further premises, and since I have offered some at least *prima facie* considerations that suggest that such an argument will be difficult to sustain, in the next section I consider a different strategy for getting at the skeptical conclusion. I go on to reject it as a path to skepticism about free will and argue that it instead suggests that there are principled reasons for rejecting skepticism about free will.

V. SKEPTICISM ABOUT SKEPTICISM ABOUT FREE WILL

A different way a skeptic might try to close the gap in the argument is to argue along these lines:

Suppose we adopt an approach modeled on Lewis's approach to defining theoretical terms—in this case, 'free will' and related notions—by isolating those properties that best correlate with our maximally consistent platitudes about free will and related notions.³⁵ If we do this, what we will find is that there is no property that sufficiently correlates with the platitudes to make talk of 'free will' sensible. Thus, we should conclude that we do not have free will.

To my knowledge, no current skeptic about free will has come to their free will skepticism on these grounds, but this should not lead us to dismiss it as a potential route to skepticism. If this route to skepticism were feasible, it would be an attractive one. First, it would permit the skeptic to acknowledge the possibility of a difference between our conceptions of free will and moral responsibility and the facts about them. The sort of skepticism generated by this approach is not one that hinges on the world being precisely like our concept. Instead, it holds that there is no property in the world that uniformly plays the same role as the conceptual role played by the predicate 'free will.' Thus, the skeptic could grant that our concept of free will may have characteristics that—if there were free will—would be distinct from the referent of 'free will.' Still, according to this skeptic, the fact of the matter is that there is no unified property, or even a nonarbitrary bundle of properties, that corresponds to a sufficiently large number of our maximally consistent platitudinous usages of 'free will,' so we should abandon the belief that we have free will.

However, one difficulty with this route to skepticism is it does little to block the possibility of revisionism about the conceptual role played by free will. That is, even if it turns out that there is no property that best fills the relevant role, we may yet find it appropriate to reorganize the conceptual role played by 'free will.' So, even if we discover that there is no property that best

corresponds to free will's conceptual role, we might discover that there are properties that correspond to significant parts of the conceptual role played by the concept of free will. Or, we might discover that there is a property that sufficiently corresponds to a slightly different conceptual role that is very similar to the overall conceptual role played by free will. In either case, we might decide that the alternative or successor conceptual role merits the name 'free will.' Thus, although the skeptic might be right about our lacking free will, given its current conceptual role, it would be a short-lived skeptical victory if we were to shift the conceptual role and the usage of the predicate 'free will' in light of that same skeptical discovery.

Whatever one thinks about the possibility of revising the conceptual role played by free will, the prior claim—that nothing corresponds well enough to the conceptual role of free will—seems implausible. Brief reflection on the conceptual role of 'free will' shows why. As previously noted, the term 'free will' is usually understood to refer to a kind of power or capacity required for moral responsibility. To be sure, there may be other conceptual roles that free will plays, and philosophers have sometimes highlighted these other roles.³⁶ However, it is clearly the case that the primary conceptual role of free will that has exercised philosophers for some many years is its role in attributions of responsibility.

Moral responsibility is itself a difficult notion to get a firm grasp on, but here too there are some platitudes to be had. Minimally, the principal role of the concept of moral responsibility is that it governs our differential treatment of people in contexts of moral significance. If I think you are morally responsible for something bad or objectionable, I treat you differently than I do if I think you are not morally responsible, or than I do if I think you are morally responsible for something good.³⁷

The details of the responsibility-characteristic attitudes and practices, while interesting and deserving of careful inquiry, are unnecessary to appreciate the more general point that the primary conceptual role of the concept of moral responsibility is to govern differential treatment of people in contexts of moral significance. Moreover, the notion of treatment is not meant to

imply that a judgment of responsibility entails physical consequences for the one judged. What I am calling “treatment” may be something as simple as the unspoken judgment that you merit praise or blame, independent of whether or not I am inclined to visit sanctions or rewards upon you.

Once we appreciate the principal conceptual role played by free will—that is, as a power or capacity relevant to our differential treatment of others—it seems somewhat less likely that there should be no such property or nonarbitrary bundle of properties in the world whose presence might sufficiently correspond to platitudinous usages of ‘free will.’ It seems highly plausible that there will be some property or nonarbitrary bundle of properties in the world that play a role that roughly corresponds to the conceptual role played by free will. The proof is to be found in the availability of a wide range of theories that identify various properties that correspond with the conceptual role of some condition that governs the differential treatment of particular agents in morally salient contexts. Whether that feature is a capacity to respond to reasons (or reasons of a particular sort, or the endorsement of one’s own motivations or values, or the capacity to guide one’s conduct on the basis of a particular quality of will, or some other thing), there are plenty of candidates for a property or nonarbitrary bundle of properties whose presence is ubiquitous enough to correspond to the conceptual role of free will—this is, after all, the business of nonskeptical theories of free will.

This is not to deny important differences between nonskeptical theories of free will, including their disagreements about how to handle particular cases. Nor is it to dismiss divergences at the margins between how we tend to think about free will and what is specified by various accounts of free will. The point is that there are a number of good accounts that specify some property that corresponds well enough to the conceptual role of free will. Thus, it is implausible to think that the skeptic is right that we will find *no* property that corresponds well enough to the conceptual role specified of free will.

At this point, the skeptic cannot object that these various property identifications (e.g., free will is the property of acting from a suitably reasons-responsive mechanism, or free will is the property of acting in accord with one's values, or free will is the property of acting in accord with one's highest-order desire, etc.) fail to identify the property that leads down the path to skepticism (such as "metaphysical ultimacy," "*causa sui*," agent-causation, etc.). Even if the skeptic is right that some such characteristic is part of our concept of free will or moral responsibility, it is no shortcoming of this approach to defining theoretical terms that it does not refer to the skeptic's favored characteristic of our concepts. The point, after all, is to identify the property (if any) that would make true a sufficiently large number of our platitudinous ascriptions of responsibility, irrespective of whether or not such a property (or bundle of properties) closely matches our conception of what we are tracking in ascriptions of responsibility. What the conceptual role skeptic would need is for there to be no property that sufficiently corresponds to the conceptual role of free will. However, as I have suggested, there are plenty of candidate properties that appear to fit that role well enough.

That there seem to be plenty of properties that correspond to the conceptual role of free will might be taken to provide a different sort of opening to the skeptic. The skeptic could argue that the abundance of candidate properties points to there being no single best property that corresponds to free will's conceptual role, which would make all ascriptions of responsibility indeterminate. Again, however, this is, at best, a pyrrhic victory. If there are a number of properties in the world that correspond to the conceptual role played by free will, we should conclude that there are a number of ways in which we can truly be said to have free will, and not that there is no free will. If we are concerned about ridding ourselves of potential semantic indeterminacy we could make something of a semantic decision, where we specify which property we aim to track, thus anchoring free will ascriptions in a unique property.

All of the foregoing considerations weigh against functional role-based skepticism about free will. However, they also suggest a positive argument for thinking that free will is not like we

conceived of it, but that we nonetheless have it. Given that the property specified by the Ramsified platitudes about free will is extremely likely to find at least one sufficiently corresponding property in the world, this should count as significant evidence in favor of the idea that we do have free will. However, for various reasons (including those previously cited concerning the vagaries of concepts susceptible to cultural accretions over time), common-sense thinking is unlikely to perfectly track whatever property or bundle of properties it is that renders true a sufficient number of the platitudes. But, as in any inquiry that is sensitive to the feedback given to us by the world, we should not be surprised about some degree of discrepancy. Rather, we should aim to make our thinking track the features of the world revealed to us as a result of our inquiry. Thus, given all the considerations we have available to us, we should suppose that we have free will, although we should also be prepared to accept that we may need to engage in some modest degree of conceptual revision if we want to line our beliefs up with the world.

One virtue of this anti-skeptical argument is that it does not face a problem encountered by one of the few other anti-skeptical arguments in the free will literature. Some have maintained that it is obvious that we are responsible, and thus, any view that is skeptical of the kind of free will that is required for moral responsibility must be false.³⁸ On the face of it, this conclusion seems unwarranted. Surely we could be wrong about whether or not we are morally responsible creatures. And, the range of attitudes and practices characteristic of our holding one another responsible may well be impossible to get rid of, as P. F. Strawson has suggested.³⁹ But our judgments that someone is responsible, and the various reactive attitudes we take toward others as a consequence do not operate in a space independent of cognitive content. As such, the beliefs or cognitive content that lurk behind our judgments and reactions may well be in error. Or, at any rate, we need some further argument or evidence about why they could not, as a matter of principle, be in error.⁴⁰ Without such an argument, there is enough space for the skeptic about free will to insert the thin edge of the skeptical wedge in this sort of anti-skeptical argument. In contrast, the anti-skeptical view I have advanced is not subject to this difficulty, for the

reference-fixing power is not uniquely settled by cognitive content. On the account I have offered, we could be in massive cognitive error (e.g., pervasively supposing that free will requires some form of libertarian agency) without this error necessarily affecting reference.

There is a final line of response that some skeptics may be tempted to deploy against the proposed anti-skeptical position. The skeptic might reply thusly: “the position that holds that we have free will, but that it is different than we imagined, is engaged in semantic hair splitting. If the anti-skeptic wants to insist that we have free will, but that free will is different than we imagined, this is not substantively different than the claim that we lack free will. In any case, we would lack the sort of thing that we have in mind when we are worried about free will, and that is reason enough to be a skeptic about free will.”

What is problematic about this line of response is that it turns skepticism about free will into a trivial or uninteresting position. On the anti-skeptical view I advanced, the notion of free will that we should embrace is the one that reflects those properties that are picked out by Ramsification of the platitudes about free will. Given the conceptual role that free will has, and given that there is very likely at least one sufficiently instantiated candidate property that corresponds reasonably well to that theoretical role, any other notion of free will will be largely uninteresting, from a philosophical perspective. Why? Any divergences from our theoretically defined notion of free will will (1) tell us little or nothing about the property in the world that is the best candidate for a truth maker for ascriptions of free will and moral responsibility and (2) tell us little or nothing about the kind of thing that defines the conceptual role that free will plays in our lives. That a theoretically defined term could turn out to be modestly revisionist does not speak against its ontological and conceptual importance any more than changes in the post-Einsteinian conception of simultaneity that is proper to physics speaks against the importance, both ontological and conceptual, of the revised notion of simultaneity. Of course, our pre-theoretical notion might persist in various parts of our linguistic community long after we have settled on a particular account of the best theoretical specification of free will. However, once we

have a successful theoretical definition on our hands, the persistence of these other notions is at best a curiosity and at worst a corrosive impediment to the spread of knowledge.

None of this is to deny that there may be cases where the proposed conceptual change is so radical as to not merit the preservation of the term across the change. However, at least in cases where the primary conceptual role of the term in question is greatly preserved, even if it is not completely preserved (as seems possible, maybe even likely, in the case of free will) it is entirely appropriate to recognize a modestly revised account of the considered notion as the notion with which we should be concerned. If so, then we should be skeptical about skeptical accounts of free will.

VI. REVISIONISM

If both libertarianism and free will skepticism are implausible, what is the incompatibilist to do? For the philosopher moved by incompatibilist intuitions, compatibilism might seem like a bitter pill to swallow, even in the face of the relatively unhappy prospects for the main varieties of incompatibilism: compatibilists have historically been too ready to dismiss libertarian concerns as conceptual confusion or spooky metaphysics invented by overly imaginative philosophers.

There is a compromise position, suggested by the arguments in the preceding section: revisionism about free will and moral responsibility. We might acknowledge that some degree of revision in our concepts (and perhaps practices) is required by a naturalistically plausible account of free will and moral responsibility. This sort of account would not require that we deny the existence of incompatibilist intuitions. On the contrary, it presupposes them and argues that we should excise them to the extent that we best can, rendering our formerly implausible concepts naturalistically plausible. Just how one might go about doing this, how it is different from various forms of compatibilism, and what all of this might entail is beyond the scope of this paper. Minimally, it requires careful reflection on the aim of a theory of responsibility, what

function is served by a system of responsibility, what forms of normatively competent agency we might be plausibly thought to have, and reflection on the psychological mechanisms involved in our responsibility-characteristic practices, attitudes, and beliefs. In generating an account of these things, we need not deny the existence of the intuitions that drive the ancient debate between compatibilists and incompatibilists. Instead, acknowledging these intuitions might keep us from being distracted from a secure path to a naturalistically plausible and normatively adequate account of free will and moral responsibility.

NOTES

Thanks to Dan Speak and Kip Werking for helpful comments on parts of this paper, to Alan “H-Bomb” Hájek for discussions about some of the ideas in this paper, and to participants of the Garden of Forking Paths weblog, <http://gfp.typepad.com>, especially Dan Speak, Rico Vitz, and Andrei Buckareff for an electronic discussion of some of the ideas in this paper.

1. I am using *skeptical incompatibilism* as a label for a family of views, including those that doubt the existence of free will and those that outright deny the existence of free will. Skepticism is typically conceived of as an epistemological view, but it seems to be used often enough in metaphysical contexts to render tolerable (although perhaps not felicitous) its use in the present context to include those views committed to the metaphysical thesis that free will does not exist.
2. On this model, some of the history of physics and many of the sciences of the mind can be understood to have developed out of philosophy along these lines.
3. P. F. Strawson, “Freedom and Resentment,” *Proceedings of the British Academy* 48 (1962). Reprinted in Gary Watson, *Free Will*, 2nd ed. (New York: Oxford University Press, 2003).
4. I have discussed the Familiar Argument and what is wrong with it in Manuel Vargas, “The Revisionist’s Guide to Responsibility,” *Philosophical Studies* 125, 3 (2005).
5. The experimental evidence for this is mixed. For evidence that common sense is compatibilist, see Eddy Nahmias et al., “Surveying Freedom: Folk Intuitions about Free Will and Moral Responsibility,” *Philosophical Psychology* (forthcoming); and Robert L. Woolfolk, John Doris, and John Darley, “Identification, Situational Constraint, and Social Cognition: Studies in the Attribution of Moral Responsibility,” *Cognition* (forthcoming). For evidence of incompatibilist elements in common sense, see Shaun Nichols, “Folk Intuitions on Free Will,” *Journal of Cognition and Culture* 6, 1 & 2 (2006); Shaun Nichols and Joshua Knobe, “Moral Responsibility and Determinism: Empirical Investigations of Folk Concepts” (unpublished manuscript). I have argued that the empirical data ultimately favors incompatibilists. See Manuel Vargas, “Building a Better Beast” (unpublished manuscript), and Manuel Vargas, “Philosophy and the Folk: On Some Implications of Experimental Work for Philosophical Debates on Free Will,” *Journal of Cognition and Culture* 6, 1 & 2 (2006).
6. This point is somewhat distinct from the point about earning the right to mock Chisholm-style agent-causal theories, discussed in J. David Velleman, “What Happens When Someone Acts?” *Mind* 101 (1992). For example, satisfaction of the standard of naturalistic plausibility does not obviously require that we can offer a reduction of all action down to its event-causal components. Perhaps the casual usage of ‘naturalism’ in

philosophical contexts such as these suggests little more than a generally pro-scientific attitude on the part of self-described naturalists in philosophy. If so, 'naturalism' is a term much like 'positive' and 'positivist' was about a century ago in Latin America—the kind of thing maybe most philosophers thought their theories should be an instance of, even if relatively few people had detailed conceptions of what that should mean.

7. See Michael C. Rea, *World without Design: The Ontological Consequences of Naturalism* (New York: Clarendon Press, 2002).
8. One important exception *may* be Peter van Inwagen, who has repudiated naturalism in print. See, for example, Peter van Inwagen, "Review of *Problems in Philosophy*," *Philosophical Review* 105, 2 (1996). It is notable, though, that he has not offered anything like a positive account of free will. For his most recent discussion of why, see Peter van Inwagen, "Free Will Remains a Mystery," in *Philosophical Perspectives*, ed. James Tomberlin (Boston: Blackwell, 2000). Also relevant to this case are my remarks below on religious libertarians.
9. For example, see Ted Honderich, *A Theory of Determinism* (New York: Oxford University Press, 1988); Saul Smilansky, *Free Will and Illusion* (New York: Clarendon Press, 2000); and Gaalen Strawson, *Freedom and Belief* (Oxford: Oxford University Press, 1986). Derk Pereboom's provocative version of hard incompatibilism is *not* committed to the view that libertarian agency is impossible, only that it is extremely unlikely to turn out true in the actual world. See Derk Pereboom, *Living without Free Will* (Cambridge: Cambridge University Press, 2001).
10. It is consistent with what I have been arguing that—at least in my experience—philosophers who are out of touch with developments in the free will literature over the past fifteen to twenty years are invariably unaware of the existence of naturalistically compatible libertarian theories.
11. Robert Kane, *The Significance of Free Will* (Oxford: Oxford University Press, 1996).
12. Laura Waddell Ekstrom, *Free Will: A Philosophical Study* (Boulder, Colo.: Westview Press, 2000).
13. Alfred Mele, *Autonomous Agents: From Self-Control to Autonomy* (New York: Oxford University Press, 1995). Note: although Mele offers an account of libertarianism, he is officially agnostic about the issue and has offered both libertarian and compatibilist accounts of what he calls "autonomous agency."
14. Randolph Clarke, "Toward a Credible Agent-Causal Account of Free Will," *Noûs* 27 (1993). Reprinted in Watson, *Free Will*, 283–98.
15. Timothy O'Connor, *Persons and Causes* (New York: Oxford University Press, 2000).
16. For an idiosyncratic statement of this sort, see Colin McGinn, *Problems in Philosophy: The Limits of Inquiry* (Cambridge, Mass.: Blackwell, 1993); and Van Inwagen, "Review of *Problems in Philosophy*." My preceding comments should make it clear why I think this sort of view is false.
17. Mark Balaguer, "Libertarianism as a Scientifically Reputable View," *Philosophical Studies* 93 (1999).
18. See Manuel Vargas, "Responsibility and the Aims of Theory: Strawson and Revisionism," *Pacific Philosophical Quarterly* 85, 2 (2004); and Vargas, "The Revisionist's Guide to Responsibility." A rare attempt to engage with methodological issues in a systematic fashion, at least as they concern the free will debate, can be found in Richard Double, *Metaphilosophy and Free Will* (New York: Oxford University Press, 1996). See also Ira Singer, "Freedom and Revision," *Southwest Philosophy Review* 18, 2 (2002).
19. For example, I suspect that Van Inwagen's willingness to give up his argument(s) for incompatibilism were determinism to be proven cannot be properly understood in absence of his background religious commitments and their implications for what moves are available to him in the debates about free will. See Peter van Inwagen, *An Essay on Free Will* (New York: Oxford University Press, 1983), 221. One might wonder why religious libertarians have apparently felt the need to adhere to a standard of naturalistic compatibility (though, again, van Inwagen may not be in this group). Possible explanations may range from the sensible (e.g., a confidence in the epistemic excellence of science and the desire to avoid conflict with it), to the more controversial (e.g., feeling bound by Rawlsian-style constraints on something like public reason) to the downright nefarious (e.g., the tyranny of analytic philosophy's sometimes evangelical atheism). And, of course, several explanations might apply in any given case.
20. For a good discussion of what is wrong with popular ways of inferring metaphysical commitments from aspects of our first-person deliberative perspective, see John Martin Fischer, "Free Will and Moral Responsibility," in *Oxford Handbook on Ethics*, ed. David Copp (Oxford: Oxford University Press, forthcoming).
21. For a recent account of a wide range of troubles that various sciences of the mind have generated for our common-sense views about the operations of our own mind, including the efficacy of introspection with respect to the will, see Daniel M. Wegner, *The Illusion of Conscious Will* (Cambridge, Mass.: MIT Press 2002), 1–28, 221–70. I recommend this book with trepidation, however. Although the Wegner book is useful for the purpose just mentioned (that is, undercutting the naïve assumption that our conscious experience tells us much about the underlying phenomena of mind and agency), it suffers from a lack of caution about the philosophical theses it advances, as has been ably demonstrated in Eddy Nahmias, "When Consciousness Matters: A Critical Review of Daniel Wegner's *The Illusion of Conscious Will*," *Philosophical Psychology* 15, 4 (2002). In a different but relevant vein, John Doris has developed a philosophically rigorous account of

- how, in particular, situations affect human behavior in ways largely invisible to us. See John Doris, *Lack of Character* (New York: Cambridge University Press, 2002), esp. 28–61.
22. Robert Kane, “Responsibility, Luck, and Chance: Reflections on Free Will and Determinism,” *Journal of Philosophy* 96, 5 (1999): 224–25.
 23. Kane, “Responsibility, Luck, and Chance,” 226.
 24. For some examples, see Henrik Walter, *Neurophilosophy of Free Will: From Libertarian Illusions to a Concept of Natural Autonomy* (Cambridge, Mass.: MIT Press, 2001), 160–63; Patricia Churchland, *Brainwise* (Cambridge, Mass.: MIT Press, 2002), 195–97; and R. Grush and Patricia Churchland, “Gaps in Perseus’s Toolings,” *Journal of Consciousness Studies* 2 (1995).
 25. Walter, *Neurophilosophy of Free Will: From Libertarian Illusions to a Concept of Natural Autonomy*, 162. More generally, see 151–63. For some of the issues that would need to get sorted out before a model of brain-based indeterminacies strikes brain scientists as plausible, see Churchland, *Brainwise*, 233–34. Although Churchland is insufficiently attentive to the differences among contemporary libertarian approaches, I take it that most of the points she makes about questions that would get raised about “uncaused decision” theories are relevant to any account of libertarian agency that is committed to there being indeterminacies in the brain.
 26. Evidence and arguments for this conclusion can be found in, for example, Alfred Mele, “Decisions, Intentions, Urges, and Free Will: Why Libet Has Not Shown What He Says He Has,” in *Topics in Contemporary Philosophy*, vol. 4: *Causation and Explanation*, ed. Joseph Klein Campbell, Michael O’Rourke, and David Shier (Cambridge, Mass.: MIT Press, forthcoming); Daniel Dennett, *Freedom Evolves* (New York: Viking, 2003), 14–21, 221–55; and Nahmias, “When Consciousness Matters.” I have pressed this point in Manuel Vargas, “Compatibilism Evolves: On Some Varieties of Dennett Worth Wanting,” *Metaphilosophy* 36, 4 (2005).
 27. Even most religious libertarians should not argue against the epistemic credentials of something like a standard of naturalistic plausibility, for though they might think that other standards need to be satisfied, standards that nonreligious philosophers might be inclined to reject, few thoughtful religious libertarians will argue against the epistemic credential implicit in the standard of naturalistic plausibility. However, to the extent that their commitment to libertarianism has grounds independent of those likely to be accepted or shared by nonreligious libertarians, the most attractive option may be something like “mysterianism”—the view that we have free will, but that it is a mystery how we have it.
 28. Despite all I have said, I am agnostic about what light, if any, these considerations cast on the rationality of libertarians as believers and the appropriateness of pursuing improbable research programs more generally. The significance of having implausible beliefs—even beliefs whose implausibility one acknowledges—for one’s status as a good belief-manager is a complicated business. Surely some libertarians will argue that considerations not relevant to the truth can play a role in the appropriateness of belief, adherence to a belief, and pursuit of an apology for the belief. If so, then perhaps libertarians are like those who believe in ESP (extrasensory perception), and those who develop libertarian theories are like researchers hoping to demonstrate the existence of ESP. There may be reasons for particular individuals to believe in the truth of ESP, and there may be benefits derived from doing research on these topics (e.g., personal satisfaction, interest, and the various unanticipated but beneficial side-effects of doing research). And, for all we know, ESP may yet turn out to be vindicated. However, from where the rest of us stand there is little reason to think ESP (or libertarianism) is true.
 29. See, for example, Double, *The Non-Reality of Free Will*; Honderich, *A Theory of Determinism*; Pereboom, *Living without Free Will*; Smilansky, *Free Will and Illusion*; and Strawson, *Freedom and Belief*.
 30. This simplified characterization of his argument is based on the version he gives in Galen Strawson, “The Impossibility of Moral Responsibility,” *Philosophical Studies* 75 (1994). Reprinted in Watson, *Free Will*, 219–20.
 31. Virtually all skeptics about free will are incompatibilists, that is, those who think that free will is incompatible with the truth of the thesis of determinism. However, one could be a compatibilist (i.e., one who believes that free will is compatible with determinism) and a skeptic about free will. Perhaps free will is incompatible with indeterminism, as Hume is often taken to have suggested, and perhaps the world really is indeterministic. However, since compatibilist skepticism has rarely, if ever, been adopted, what follows largely ignores this possibility. On the characterization of incompatibilism I am using, even figures such as Honderich and Smilansky, both of whom resist the label ‘incompatibilist,’ turn out to be incompatibilists, because they think that free will (in the sense I specified above) is incompatible with determinism. It is also worth noting that many skeptics about free will think that we lack it, even if determinism is *not* true (e.g., Pereboom and Strawson).
 32. Depending on the actual concept of marriage, recent public debates about gay marriage may yield some amount of conceptual revision in the concept of marriage. Even so, this would only be the latest in a long series of revisions in how we understand marriage. I take that the contemporary Western concept of marriage is one that is considerably different from the older understanding of marriage as a property exchange between the father and husband of the bride.

33. We should be open about what constitutes post-Quinean conceptual analysis, but it is presumably something akin to the “concept mapping” described in P. F. Strawson, *Analysis and Metaphysics* (New York: Oxford University Press, 1992).
34. It is no reply to the foregoing to respond with “whatever the notion of freedom is that we would have post-conceptual change would not be freedom in the sense required for ultimate responsibility.” The reason is that the sense of freedom under discussion precisely is that sense of freedom required for responsibility in the deep or ultimate sense.
35. See David Lewis, “How to Define Theoretical Terms,” in *Philosophical Papers* (New York: Oxford University Press, 1983). For a related proposal that focuses specifically on moral terms (something ‘free will’ conceivably is), see Frank Jackson and Philip Pettit, “Moral Functionalism and Moral Motivation,” *Philosophical Quarterly* 45, 178 (1995).
36. See, for example, Robert Nozick, *Philosophical Explanations* (Oxford: Clarendon Press, 1981), 291–93, and the list of issues canvassed in Kane, *The Significance of Free Will*, 80.
37. This picture is compatible with views that hold that there is some fact of the matter about whether or not someone is responsible, as well as views that reject the possibility of there being such facts.
38. See Van Inwagen, *An Essay on Free Will*, 162, 188.
39. Strawson, “Freedom and Resentment.”
40. An argument of this sort has been made in Susan Wolf, “The Importance of Free Will,” *Mind* 90 (1981). I discuss this sort of argument in the context of Strawsonian approaches to free will in Vargas, “Responsibility and the Aims of Theory: Strawson and Revisionism.”