

FREE WILL AND TIME TRAVEL

Neal A. Tognazzini

When thinking about potential threats to free will, people tend to think big. If there is an essentially omniscient and everlasting God, for example, then what God believed millions of years ago about the future entails every true proposition about what we're up to now, and since it's not up to any of us what God believed millions of years ago, what we're up to now isn't up to any of us, either (Pike 1965). Alternatively, if causal determinism is true, then the laws that govern our world, together with facts about how the world was millions of years ago, entail every true proposition about what we're up to now, and since it's not up to any of us what the laws are like or what the world was like millions of years ago, what we're up to now isn't up to any of us, either (van Inwagen 1983). Finally,—and I mention one last example here because we'll return to it later— even if there is no God, presumably there were, millions of years ago, true future-tensed statements about what we would be up to millions of years later, and those true future-tensed statements entail(ed) all the present-tense propositions about what we're up to now. Again, since it's not up to us what was true millions of years ago, what we're up to now isn't up to us, either (Taylor 1963).

These three threats to free will—the threat from logical determinism, causal determinism, and God's foreknowledge—are all *global* threats, because they are general theses about the nature of reality whose alleged consequences would be disastrous for every creature whatsoever. But there are *local* threats to free will that are worth considering, too, puzzles about relatively well-circumscribed portions of reality that appear to threaten this or that person's free will. In this chapter, I intend to explore one of these more modest threats, one that arises from endorsing the possibility of time travel to the past.

Most discussions of free will and time travel occur in contexts where what's at issue is whether time travel is metaphysically possible. In those contexts, the worry is that if it were possible for someone to travel to the past, then contradictions could be true. For example, the time traveler would be able to kill his grandfather (because he would train for the occasion, he would have a gun, he would know where to find his grandfather, etc.), and also would not be able to kill his grandfather (because if he were to kill his grandfather, then he himself wouldn't exist to travel back in time in the first place). There seems to be general agreement, at least in the last 40 years, that these sorts of worries do not show that time travel to the past is metaphysically impossible (Lewis 1976; R. Wasserman, unpublished manuscript, "*The Paradoxes of Time Travel*"). I agree,

but I am not concerned with whether time travel is metaphysically possible. Instead, I plan to take for granted that it is, and then ask what it would mean for the free will of the time traveler. It's a common thought that time travelers to the past would somehow be constrained to do only those things that they in fact did do, that they wouldn't be free to do anything that didn't already happen. That is, it's a common thought that we should be *incompatibilists* about free will and time travel to the past. (Note that this is a different sense of the term 'incompatibilist' than is typical in debates about free will and, say, determinism.) But is this common thought right?

In what follows, I'll suggest that the answer to this question is complicated. On the traditional way of thinking about free will, I think the incompatibilist about time travel and free will wins the day. However, I'll also consider a residual worry with the incompatibilist conclusion, one that contains the seeds of a new—and, I think, more promising—way of thinking about free will, about what it is for an action to be up to someone.

Models of Time Travel

To begin, let's distinguish between three different models for travelling to the past: the single timeline model, the branching timeline model, and the hypertime model. We'll be concerned with the single timeline model, but I'll briefly describe the others for the sake of contrast. To understand the models, I find it useful to assume that spacetime is a four-dimensional block, which contains every spatiotemporal thing that has ever, does now, or will ever exist. (I say that I find this assumption useful, but let me hasten to add that it is dispensable [Keller and Nelson 2001; van Inwagen 2010].) You and I are in the block, but so are Abraham Lincoln and the people whose birth year is 2115. Now suppose that Harry is born in 2100, enters a time machine in the year 2120 and travels back 100 years to 2020, shakes his grandfather's hand, and then goes back to the future, arriving one minute after he left.

On the single timeline model, if you imagine scanning the four-dimensional spacetime block from left to right, starting at the year 2019, here's what you would see (you'll have to suppress a couple of spatial dimensions in your mind's eye): in the year 2020, a 20-year-old Harry appears, with his time machine, seemingly out of nowhere. He exits his time machine, finds his grandfather, and shakes his hand. Then he re-enters the time machine and disappears. Harry's grandfather then grows up, gets married, and has children (one of whom is Harry's mother), and then little baby Harry himself is born in 2100, grows into a time travel enthusiast who finally enters a time machine in the year 2120, and then disappears with his time machine. After a minute passes (you're still scanning the block left-to-right), Harry and his time machine reappear (both are more than a minute older), Harry exits, and he goes about his life. On this model, Harry does not cause anything to happen in the past that didn't (already) happen in the past. It may sound odd, but as Harry departs, we can truly say of him that he is about to experience an event that has already taken place. (Does that mean that the event is both future and past? Yes and no. If we're talking about *external* time, the event Harry is about to experience is only past. But if we're talking about Harry's *personal* time [the ordering of his life experiences from younger to older], then the event is only future [see Lewis 1976].)

The branching timeline model, on the other hand, promises to open up the possibility that what the time traveler does on his journey into the past is different than what

happened ‘the first time around.’ (On the single timeline model, there is only one ‘time around.’) To picture things on the branching timeline model, we can again imagine a four-dimensional block universe, except now it has a branching structure. There is a single trunk from the beginning of time until the year 2020, but at 2020 the trunk sports a branch, so that the four-dimensional block universe at that point looks like a ‘y’ that has been laid on its side. The trunk, as it continues past the sprouted branch, contains the ‘first time around,’ where Harry’s grandfather lives out his life *sans* handshake. But the sprouted branch contains the ‘second time around,’ where Harry arrives in his time machine and shakes his grandfather’s hand. (I’ll ignore the very good question of whether this story really counts as time travel.)

Finally, consider the hypertime model (Goddu 2003; van Inwagen 2010; Hudson 2014). The idea here is that in addition to the sequence of moments *within* the four-dimensional spacetime block itself, there is also a sequence of hypermoments (moments in a second-dimension of time) through which the very shape of the spacetime block might morph and change. Picture it this way: you are outside of the spacetime block (with a God’s-eye perspective on all of spatiotemporal reality—past, present, and future) and you have a reel-to-reel video camera. As hypertime passes (for you, who are outside of the spatiotemporal block), you are capturing what the block looks like from hypermoment to hypermoment. (Each hypermoment is represented by a frame on your filmstrip.) If the spacetime block hypernever changes its shape, then you are filming a rather boring movie—just like staring at a loaf of bread that never rots (so, perhaps like staring at a Twinkie). But once we have the idea of hypertime, we can use it to imagine a shapeshifting block. Truths about what is (simply) *past* are always about what the leftward sections of the spacetime block contains, but truths about what is *hyperpast* are about what the filled-up frames on your filmstrip represent the spacetime block as (hyper)having contained at those frames. To imagine time travel on this model, all we have to do is to imagine that what the spacetime block contains at one hypermoment is different from what it contains at the next hypermoment. At the first hypermoment, the block contains a grandfather who never got to meet his grandson. But at the next hypermoment, the spacetime block contains not only the 20-year-old Harry who is spooling up his time machine in the year 2120, but also a slightly older Harry who is shaking hands with his grandfather in the year 2020. (The contents of the 2020 slice of the spacetime block have changed from what they hyperwere.)

The particular (local) incompatibilist argument I’ll consider arises most naturally (or, perhaps, only) on the single timeline model of time travel. I’ve sketched the others only for the sake of contrast, so that the reader doesn’t automatically begin thinking in terms of the *Back to the Future* movie franchise, with its fading photographs and changing of the past. The worry for free will that arises in the context of single timeline time travel arises in part because of the natural thought that we *can’t* change the past.

A Worrisome Time Travel Story

To home in on the worry I have in mind, let’s talk about another Harry, this one from *Harry Potter and the Prisoner of Azkaban* (Rowling 2001). Here’s the background you need: Harry is a wizard who knows how to cast a spell to ward off dark creatures called Dementors (the spell is called a Patronus), and he has a friend named Hermione who has a device called a timeturner, which looks like a charm necklace but which allows its

wearer(s) to travel through time. And here's how our story goes (told from the perspective of Harry's personal time): at t_1 Harry's life is threatened by several Dementors, but at t_2 he is saved because someone lurking in the shadows casts a Patronus spell at the Dementors. Hours later, at t_3 , he and Hermione use a timeturner to travel back to t_1 , where Harry sees himself (his younger self) threatened by the Dementors. At t_2 , Harry casts a Patronus spell at the Dementors, thus saving his own life. Just before t_3 , Harry and Hermione use the timeturner to travel back to the future, and they arrive just in time to see themselves (their younger selves) disappearing at t_3 to begin their adventures in the past.

This is a single timeline story of time travel to the past. Harry saves his own life at t_2 , and because he does, it is true at t_3 that he saved his own life, even though the younger version of him at t_3 (the one about to head back in time) doesn't realize that this is what happened (he knows *someone* saved his life, but he doesn't know it was/will be *him*). As the story unfolds in the book, the realization that the person in the shadows was (is) himself is liberating. As Harry puts it (Rowling 2001: 412), "I knew I could do it [i.e. cast the Patronus spell] this time because I'd already done it!" In other words, since he knew that he *would* save himself (talking from his perspective before casting the spell, but with the knowledge that he had lived long enough to travel back in time), he knew that he was *able* to save himself. This realization is, in fact, what *enabled* him (perhaps it's what allowed him to summon the necessary concentration). But by the same token, we might wonder whether saving his own life was a free action of his—whether, given the peculiar circumstances, it really was up to him whether he cast the Patronus. After all, had he failed to cast the Patronus, he wouldn't have survived the Dementor attack, and hence wouldn't exist to even be faced with the choice of whether to save himself.

Lest we think that this sort of worry only arises for agents who are in the midst of a time-travelling adventure, it's worth pointing out that there is another question that arises from the story as we have told it, namely: was Harry's action of travelling back in time at t_3 a free action? At t_3 , (younger) Harry has not yet experienced the thrill of traveling through time. But the fact that he will soon arrive in the past and save his own life leads to uncomfortable questions about his freedom just the same. After all, had he failed to travel back in time, nobody would have been in the shadows ready to save him from the Dementor attack (we can stipulate that Hermione, even if she were to have traveled to the past alone, wouldn't have been able to cast a strong enough Patronus), and so he wouldn't exist to even be faced with the choice of whether to go back in time. The puzzle to be considered in what follows isn't only about the free will of *time travelers while they are time traveling*; rather, it's about free will, given the possibility of time travel.

On the traditional model of thinking about free will, whether an action is up to someone (i.e., whether it is free) is a matter of whether the person is both *able* to perform the action and also *able* to refrain from performing it (van Inwagen 1983). Adopting that model for the moment, our questions about Harry become:

Q1: At t_2 , is (older) Harry able to refrain from casting the Patronus?

Q2: At t_3 , is (younger) Harry able to refrain from going back in time?

It seems to me that these questions will have the same answer, so for ease of exposition I will focus only on Q2.

Time Travel and Fatalism

In the previous section, I sketched the argument for the conclusion that Harry isn't able to refrain from travelling back in time at t_3 . But the sketch I gave relied on some unstated premises that are best brought into the open. So consider how we might regiment the argument:

1. (Older) Harry casts the Patronus at t_2 .
2. Necessarily, (older) Harry casts the Patronus at t_2 only if he travels back in time at t_3 (where t_2 is externally past, relative to t_3).
3. So, if Harry hadn't traveled back in time at t_3 , then some fact about the past would have been false.
4. If, in order for S to do A, some fact about the past would have to be false, then S isn't able to do A.
5. Therefore, Harry isn't able to refrain from travelling back in time at t_3 .

The crucial ingredient of this argument is premise (4), which captures the idea that the past is fixed—that is, the past is not something that anyone is able to do anything about anymore. (No use crying over spilled milk, and all that.) This premise—we might call it The Principle of the Fixity of the Past (FP)—plays a crucial role in the global incompatibilist arguments we canvassed at the beginning of this chapter, so it shouldn't be surprising that it shows up in this context too. It's precisely because God's beliefs are *in the past* that they seem to cause trouble for free will, and it's precisely because the thesis of causal determinism implies that our actions are entailed *by the past* and *the laws* (which also seemed fixed) that it seems to cause trouble for free will. (Note that I'm using the word 'entail' in a broadened sense throughout this chapter: something of one ontological category entails something of a different ontological category, in this sense, if it's not metaphysically possible for the first thing to obtain [or exist, or happen, or be true, etc.] without the second thing obtaining [or existing, or happening, or being true, etc.].) Likewise, what's causing trouble for Harry here is that there are facts about the past that seem to entail what he does. And things are even worse for Harry, because the fact that he's a time traveler means that he has *two* pasts—the external past and his personal past—and both contain events that entail his actions at future times (times in the external future and times in his personal future).

So how can a compatibilist about time travel and free will respond? David Lewis is perhaps the Platonic Ideal of a compatibilist, and what he would say is that this argument is “a bit of fatalist trickery” (Lewis 1976). His remarks explaining that verdict are worth quoting in full:

Fatalists—the best of them—are philosophers who take facts we count as irrelevant in saying what someone can do, disguise them somehow as facts of a different sort that we count as relevant, and thereby argue that we can do less than we think . . . I am not going to vote Republican next fall. The fatalist argues that, strange to say, I not only won't but can't; for my voting Republican is not compossible with the fact that it was true already [in the past] that I was not going to vote Republican . . . My rejoinder is that this is a fact, sure enough; however, it is an irrelevant fact about the future masquerading as a relevant fact about the past, and so should be left out of account in saying what, in any ordinary sense,

I can do. We are unlikely to be fooled by the fatalist's methods of disguise in this case, or other ordinary cases. But in cases of time travel, precognition, or the like, we're on less familiar ground, so it may take less of a disguise to fool us. Also, new methods of disguise are available, thanks to the device of personal time.
(Lewis 1976: 151)

The fatalist is someone who worries that since future-tensed propositions about our actions were true millions of years ago, and since the past truth of those propositions entails the present truth about what we are up to now, we don't have free will with respect to anything that we ever do. What Lewis is saying in the above quotation, however, is that some facts that *appear* to be about the past aren't *really* about the past; they are simply masquerading. In particular, *all* of the facts that the fatalist points to—facts about which future-tensed propositions were true millions of years ago—are not genuinely about the past. Since they aren't, Lewis says that we shouldn't take them into consideration when trying to determine what I'm able to do in the here and now.

Another way of saying this is to say that since those facts aren't genuinely about the past, they aren't covered by The Principle of the Fixity of the Past. More carefully stated, (FP) only applies to facts that are genuinely about the past, and not every fact that was *true* in the past is really *about* the past. In the language of Ockhamism (Pike 1965; Fischer 1989), there's a distinction to be drawn between *soft* and *hard* facts—facts that are temporally relational (soft) versus facts that are temporally intrinsic (hard)—and it's only the *hard* facts that are plausibly thought to be fixed. So (FP) should really be more carefully stated as the Principle of the Fixity of the Hard Past:

(FHP) If, in order for S to do A, some hard fact about the past would have to be false, then S is not able to do A.

But now if *this* principle appears as premise (4) in the above argument, then the argument turns out to be invalid. The 'past fact' that entails that Harry travels back in time at t_3 —namely, the fact that (older) Harry casts a Patronus at t_2 —is a temporally relational (soft) fact, since whether it obtains depends crucially on what happens in the future. In particular, whether Harry casts the Patronus at t_2 depends on whether he goes back in time at t_3 . So, although it is a fact that obtains at t_2 , it's not wholly *about* t_2 , and hence isn't one of the facts that we need to hold fixed when deciding what Harry is able to do at t_2 .

All this is a way of saying: if Ockhamism seems like a good solution to the problem of logical fatalism—and to many people, it does—then it's also a good solution to the puzzle about time travel and free will. In cases of time travel, past and future go all topsy-turvy, but if we carefully separate personal past from external past, and soft past from hard past, then we should be able to sort things out properly, and we'll end up compatibilists. What's tricky is just that some of the time traveler's personal past lies in the external future (and so need not be held fixed, even though it's past in some sense), and some of the time traveler's personal future lies in the external past (and so may be a soft fact that need not be held fixed, even though it's in the external past).

Incompatibilism Revived

Pointing out that sometimes facts about the future masquerade as facts about the past is helpful in responding to fatalism, and it's also helpful in giving a first response to the

incompatibilist worry about time travel. But there seems to be something *else* that's worrisome about time travel, even once we do our best to avoid being fooled by fatalist trickery. We can bring out this additional worry by comparing the following two questions:

1. Suppose it was true a million years ago that I will have a cup of coffee tomorrow morning. If I'm nevertheless able to refrain, then it must be possible for me to refrain. What would the world be like if I were to refrain?
2. Suppose that (older Harry) did cast the Patronus at t_2 . If Harry is nevertheless able to refrain from going back in time at t_3 , then it must be possible for him to refrain. What would the world be like if he *were* to refrain?

The first question seems like it has a relatively straightforward answer: if I were to refrain from having that cup of coffee, then something that was in fact true would be false. That thought is perhaps initially worrying, but once we make the distinction between hard and soft facts, the worry dissolves. If I were to refrain from having that cup of coffee, only some *soft* fact about the past would be false.

But the answer to the second question, while equally straightforward, seems much more worrisome. If Harry were to refrain from going back in time at t_3 , then he would have died at an earlier time, and hence wouldn't have existed at t_3 . But if he hadn't existed at t_3 , then he wouldn't have been around to *do* anything at t_3 , and hence wouldn't have been around to refrain from going back in time at t_3 . Since Harry's very existence at t_3 seems to rely on something he is about to do in his personal future, this situation seems more puzzling than the one above. But what exactly is the problem?

To see what it is, consider the following two principles, neither of which seems to rely on the thought that the past is fixed, but both of which seem to generate the incompatibilist worry nonetheless (Vihvelin 1996; Spencer 2013).

(KV) S is able to do A only if, had S tried to do A, S would or at least might have succeeded.

(JS) S is able to refrain from doing A only if, had S not done A, S would have done something else instead.

Both principles specify necessary conditions on someone's being able to do something, and in each case the necessary condition is that certain counterfactual should come out true. According to (KV), roughly, in order for S to be able to do A, there must be some relatively nearby world in which S tries to do A and succeeds. The thought here is that if S doesn't do A but nevertheless *could have*, then a world in which S does A must have been *accessible* to her. And although the notion of accessibility is admittedly obscure, it seems clear enough that if she wouldn't succeed in doing A, even if she had tried, then she lacks access to a world in which she does A—that is, she wasn't able to do A.

But you might worry about putting too much weight on the notion of *trying* (Vranas 2010), in which case principle (JS) can do the same work. Imagine we're in a world in which S does A, and we're wondering whether S was able to do otherwise. The thought behind principle (JS) is that it's not enough to find nearby worlds in which it's false that S does A, because there are two ways that it can be false that S does A. First, and most

naturally, S might do something else. But another way it could turn out false that S does A is if S doesn't exist, and so isn't around to do anything at all. What principle (JS) says is that those worlds without S don't count when it comes to determining whether S was able to do otherwise. If there isn't an accessible world where S is doing something other than A, then S isn't able to do anything other than A.

Focusing on principle (JS), let's now return to Harry. In fact, Harry traveled back in time at t_3 , but if he hadn't, what would he have been doing instead? It seems like we have to say that there's *nothing* he would have been doing instead, because the nearby worlds in which he doesn't go back in time at t_3 are worlds in which younger Harry dies at t_2 and hence are worlds in which there is no older Harry at t_3 at all. But what this means is that if principle (JS) is true, then Harry is not able to refrain from traveling back in time at t_3 . In the form of an argument:

1. Harry is able to refrain from traveling back in time at t_3 only if, had Harry not traveled back in time at t_3 , he would have done something else instead.
2. But if Harry had not traveled back in time at t_3 , then he would not have existed at t_3 (and so wouldn't have been doing something else).
3. It's possible that Harry doesn't travel back in time at t_3 (since it's possible that Harry never existed at all). (On why this premise is needed, see Spencer 2013.)
4. So, it's not the case that, had Harry not traveled back in time at t_3 , then he would have done something else instead.
5. Therefore, Harry is not able to refrain from traveling back in time at t_3 .

If this argument relies on any sort of fatalist trickery, it's at least not obvious what the trick is. Instead, what it seems to rely on is a rather plausible principle that specifies a necessary condition on someone's being able to refrain from doing something, along with a couple of other plausible counterfactuals. Of course, it's controversial how to evaluate counterfactuals like those that appear in the premises of this argument (Lewis 1973), but there doesn't seem to be anything special about these.

So, Lewis is right that we don't ordinarily hold soft facts fixed when deciding what someone is able to do. Nevertheless, claims about what people are able to do seem to require the truth of some rather mundane counterfactuals, and those counterfactuals come out false in time travel scenarios. The incompatibilist challenge can be revived even once it is liberated of its fatalist attire.

A Residual Worry

One bothersome aspect of the incompatibilist conclusion is that it seems to commit us to the existence of "strange shackles" on the time traveler. That's Ted Sider's phrase, and he continues: "once the inability of the time traveler to kill her former self is admitted, one wonders what prevents her from doing so" (Sider 2002: 122). Surely we don't have to suppose that there is some "guardian of logic" (2002: 132) that disables the time traveler from doing things we would ordinarily say non-time-travelers could do.

One way we might respond to this worry is to point out that the arguments that have been given for the conclusion that time travelers are able to kill their grandfathers all fail (R. Wasserman, unpublished manuscript, "*The Paradoxes of Time Travel*"). Another way would be to say that the notion of freedom is, after all, a modal notion, so it shouldn't be at all surprising if it turns out that truths about what would happen if things


were different render an agent unfree, even if there seems to be no *mechanism* by which his lack of freedom is secured (but see Byerly 2014).

But I actually think that there is a way of refining the “strange shackles” objection to give it more force. And the way in which we refine it will, in fact, lead us to an alternative model of what it is for an action to be *up to* an agent—that is, to a model of free will that is different from the model according to which free will is a matter of being both able to perform an action and able to refrain. On this new model, compatibilism about time travel and free will can (perhaps) be resurrected.

Entailment vs. Sourcehood

The big incompatibilist worries mentioned at the outset all share a structure: in each case, there is something that’s not up to anyone but that seems to be the source of what everyone does. But the worries as I outlined them above invoked the notion of *entailment*, and we might wonder whether entailment is a legitimate and worrying sense of *sourcehood*. That is, I’ve presupposed that when p entails q , there’s nothing wrong with saying that p is the *source* of q . But is that right?

To see why we might doubt it, consider again the distinction between hard and soft facts. The response to the logical fatalist, and to the argument for time travel incompatibilism that we considered at first, is that the facts at issue are actually only *soft* facts about the past, and thus that the Principle of the Fixity of the Past, once properly understood as applying only to hard facts about the past, doesn’t show us that they are fixed. More work would need to be done in order to show that those facts *aren’t* fixed, of course, but here’s one observation that takes us some distance toward that conclusion: those past facts about what I will go on to do are true, if they are true at all, at least partly in virtue of what I do in the future. In other words, those past facts aren’t the *source* of what I do. If anything, the explanatory relationship is the other way around.

What this observation shows us, if it’s right, is that some fact about the past might entail what I do without being the source of what I do. (Similar considerations will be at play in the debate about causal determinism and free will if we are inclined to adopt a Humean view of the laws of nature [Beebe and Mele 2002].) So, it does  follow from the fact that my actions are entailed by some fact about the past that fact about the past isn’t at least partly up to me. Although I don’t have the space here to argue for the point, I’m inclined to think that what’s central to a proper understanding of free will aren’t facts about whether an agent is able to do otherwise, but rather facts about whether the allegedly free actions occur in virtue of relevant facts about me and my mental states. What’s worrisome about the thesis of causal determinism (if we don’t go Humean about the laws) is that if it is true, my actions seem to occur in virtue of stuff that isn’t me. And that’s precisely why God’s foreknowledge has seemed less worrisome to many, because his beliefs—even if he held them millions of years ago—are what they are at least partly in virtue of what I do in the future (Todd 2013; Tognazzini 2015).

But now return to the “strange shackles” worry. I submit that the motivation behind this worry is the thought that although the time travel story might show us that a fact about the past (namely, Harry’s casting the Patronus at t_2) entails that Harry will travel back in time at t_3 , nothing about the time travel story seems to suggest that there’s anything *beyond Harry himself* that is the source of Harry’s decision to travel back in time at t_3 . The reason that literal shackles can disable us isn’t simply the fact that our being in them entails that we won’t perform certain actions. Rather, it’s that our not

performing those actions is *explained* by the fact that we are in shackles, together with the fact that the shackles are something outside of our own agency (cf. Frankfurt 1969). So if you find it bizarre that Harry could somehow be rendered unfree merely by the fact that his future self did something in the past which entails that he'll perform a certain action in the present, then perhaps what's worrying you is that it's *Harry's actions themselves* that are doing the entailing, and surely Harry can't be 'shackled' by something that he himself does. (Well, I suppose that Ulysses does tie himself to the mast. But the point is that shackles are worrisome for free will only when they aren't self-imposed and nevertheless do explanatory work.)

So, what I'd like to suggest in closing is that the "strange shackles" objection, once properly understood, actually points us toward a different conception of free will, one according to which what matters is not whether an agent is able to do otherwise, but whether the action occurs *in virtue of* some relevant facts about the agent and his mental states. It's puzzling to conclude that Harry lacks free will, precisely because it looks like *he* is involved in everything that allegedly constrains him. And how could he be his own shackles? If truths about abilities are governed by principles such as (JS), and questions about free will are questions about abilities, then Harry can be shackled by actions that lie in his personal future. To revive compatibilism, we could always try to find some way to reject (JS). But what I'm suggesting is that maybe a better way to go is to reject the link between free will and the ability to do otherwise. Perhaps my actions can be *up to me* even though I'm not able to refrain from performing them. Unfortunately, though, I don't have space to do any more than simply leave that as a suggestion.

Acknowledgments

For conversations that have shaped my thinking about these issues, thanks very much to John Martin Fischer, Kelly McCormick, Joshua Spencer, Patrick Todd, Ryan Wasserman, and the students in my 2014 freshman seminar (Metaphysics for Muggles) at The College of William & Mary.

References

- Beebe, H. and Mele, A. (2002) "Humean Compatibilism," *Mind* 111: 201–23.
- Byerly, T.R. (2014) *The Mechanics of Divine Foreknowledge and Providence*. New York: Bloomsbury.
- Fischer, J.M. (ed) (1989) *God, Foreknowledge, and Freedom*. Palo Alto: Stanford University Press.
- Frankfurt, H. (1969) "Alternate Possibilities and Moral Responsibility," *The Journal of Philosophy* 66: 829–39.
- Goddu, G.C. (2003) "Time Travel and Changing the Past (Or How to Kill Yourself and Live to Tell the Tale)," *Ratio* 16: 16–32.
- Hudson, H. (2014) *Hypertime and the Fall*. Oxford: Oxford University Press.
- Keller, S. and Nelson, M. (2001) "Presentists Should Believe in Time-Travel," *The Australasian Journal of Philosophy* 79: 333–45.
- Lewis, D. (1973) *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1976) "The Paradoxes of Time Travel," *American Philosophical Quarterly* 13: 145–52.
- Pike, N. (1965) "Divine Omniscience and Voluntary Action," *The Philosophical Review* 74: 27–46.
- Rowling, J.K. (1999) *Harry Potter and the Prisoner of Azkaban*. New York: Scholastic.
- Sider, T. (2002) "Time Travel, Coincidences, and Counterfactuals," *Philosophical Studies* 110: 115–38.
- Spencer, J. (2013) "What Time Travelers Cannot Not Do (But Are Responsible for Anyway)," *Philosophical Studies* 166: 149–62.
- Taylor, R. (1963) "Fatalism," *The Philosophical Review* 71: 56–66.
- Todd, P. (2013) "Soft Facts and Ontological Dependence," *Philosophical Studies* 164: 829–44.

NEAL A. TOGNAZZINI

- Tognazzini, N. (2015) "Grounding the Luck Objection," *The Australasian Journal of Philosophy* 93: 127–38.
van Inwagen, P. (1983) *An Essay on Free Will*. Oxford: Clarendon Press.
van Inwagen, P. (2010) "Changing the Past," *Oxford Studies in Metaphysics* 5: 3–40.
Vihvelin, K. (1996) "What Time Travelers Cannot Do," *Philosophical Studies* 81: 315–30.
Vranas, P. (2010) "What Time Travelers May Be Able to Do," *Philosophical Studies* 150: 115–21.

Further Reading

- Arntzenius, F. (2006) "Time Travel: Double Your Fun," *Philosophy Compass* 1: 599–616. (An entertaining and accessible introduction to time travel and some of its worries).
Carroll, J. (2010) "Context, Conditionals, Fatalism, Time Travel, and Freedom," in J. Campbell, M. O'Rourke, and H. Silverstein (eds), *Time and Identity*. Cambridge: MIT Press, pp. 79–93. (These three essays are contributions to the discussion of Kadri Vihvelin's thought-provoking article [see above]).
Fischer, J.M. and Todd, P. (eds) (2015) *Freedom, Foreknowledge, and Fatalism*. New York: Oxford University Press. (A state-of-the-art introduction to, and a collection of essays about, the issues of foreknowledge and fatalism).
Kiourti, I. (2008) "Killing Baby Suzy," *Philosophical Studies* 139: 343–52.
Vranas, P. (2009) "Can I Kill My Younger Self? Time Travel and the Retrosuicide Paradox," *Pacific Philosophical Quarterly* 90: 520–34.

Related Topics

- Classical Compatibilism
Logical Fatalism
Free Will and Theological Fatalism

Taylor & Francis
Not for distribution