

Freedom, Foreknowledge, and the Principle of Alternate Possibilities

KADRI VIHVELIN
University of Southern California
Los Angeles, CA
USA 90089-0451

For us there is only the trying. The rest is not our business.
T.S. Eliot, *Four Quartets*

Few arguments in contemporary philosophy have had more influence than Harry Frankfurt's 'Alternate Possibilities and Moral Responsibility.'¹ In that famous paper Frankfurt noted that all parties to the traditional debate about the compatibility of free will and moral responsibility with determinism had subscribed to a common assumption. They had assumed the truth of something Frankfurt called 'the Principle of Alternate Possibilities,' which he expressed as follows:

(PAP) A person is morally responsible for what he has done only if he could have done otherwise.

In the traditional debate incompatibilists argued that if determinism is true, then no one can ever do otherwise, while compatibilists argued that there is a morally relevant sense in which even a deterministic agent can

1 Harry Frankfurt, 'Alternate Possibilities and Moral Responsibility,' *Journal of Philosophy* 66 (1969) 829-39, reprinted in Frankfurt's *The Importance of What We Care About: Philosophical Essays* (New York: Cambridge University Press 1988), 1-10. Page references are to the reprinted version.

do otherwise.² Frankfurt proposed to show that **PAP** is false, thereby undercutting the traditional debate.

Whatever we think of Frankfurt's argument, we must agree that Frankfurt succeeded in changing the way philosophers think about these issues. A principle which was once almost universally accepted as a commonsense truism or even an a priori truth is now widely regarded as false, or, at the very least, highly controversial. Among compatibilists, I think it fair to say, the received view is that **PAP** is false and that therefore the traditional free will/determinism debate is irrelevant to questions of moral responsibility. More surprisingly, some libertarians have also been convinced that someone might be responsible even though he is unable to do otherwise. The upshot is that there is now a literature devoted to the challenging task of providing a new rationale for thinking that determinism deprives us of freedom and responsibility. At the same time, there are philosophers who strenuously deny that Frankfurt has succeeded in showing that **PAP** is false.

It's difficult to explain to someone not working in this area just how peculiar the situation is. On the one hand, Frankfurt stories, as they have come to be called, have had an impact in free will circles that is comparable to the impact of Gettier stories in epistemology. On the other hand, after over thirty years of debate and discussion, it is *still* controversial whether Frankfurt or any of his followers have succeeded in providing a genuine counterexample to **PAP**.

If Frankfurt's aim was to convince libertarians that even if determinism renders us unable to do otherwise, it does not undermine responsibility, he has failed. If his aim was to make it easier to defend compatibilism, he has failed. And if his aim was to bypass questions about the truth-conditions of 'can do otherwise' claims, he has also failed, for the debate that has arisen in the wake of his original thought

2 At one time, discussions of the free will problem began with G.E. Moore's observation that if determinism is true, then there is one sense of 'can' in which no one can ever do otherwise: it is inconsistent with the past and the laws of nature that anyone does other than what they actually do. On the other hand, if by 'S can do otherwise,' we mean something like 'if S chose to do otherwise, he would' or 'if S had what he took to be good reasons for doing otherwise, he would,' then the truth of determinism is consistent with its being sometimes true that agents can do otherwise. Until the publication of Frankfurt's article, it had been assumed that the question of the compatibility of freedom and moral responsibility with determinism turned on the question of whether any conditional (or, perhaps, dispositional) analysis of 'can' succeeds in capturing what we ordinarily mean when we say, in contexts relevant to moral obligation and responsibility, 'he could have done otherwise' (G.E. Moore, *Ethics* [New York: Oxford University Press 1911], chap. 6).

experiment is now mired deep in the very metaphysical questions he sought to avoid.

But in a way Frankfurt has been successful. The main debate in current free will/determinism literature is no longer about whether determinism is compatible with the ability to do otherwise; it is about whether Frankfurt stories succeed in showing that the *inability* to do otherwise is compatible with moral responsibility.³

It is my view that this literature is a philosophical dead end. Although I am a compatibilist, I think that Frankfurt's strategy for defending compatibilism is a bad one. If we begin with the commonsense view that someone is morally responsible only if she could have done otherwise, then Frankfurt stories will not and *should not* change our minds. If we are persuaded by Frankfurt, it is because we have been taken in by a bad argument.

Frankfurt is not usually thought of as providing an argument. The literature talks about Frankfurt stories as if they were, like Gettier stories, ingenious counterexamples to a once widely accepted philosophical thesis. So let's begin by seeing how such stories are supposed to change our minds about what's necessary for moral responsibility.

3 The literature is enormous and shows no signs of slowing down. Here's a partial list: Donald Davidson, 'Freedom to Act,' in his *Essays on Actions and Events* (New York: Oxford University Press 1973); Robert Heinaman, 'Incompatibilism without the Principle of Alternative Possibilities,' *Australasian Journal of Philosophy* 64 (1986) 266-76; Ishtiyaque Haji, 'A Riddle Concerning Omissions,' *Canadian Journal of Philosophy* 22 (1992) 485-502; Timothy O'Connor, 'Alternative Possibilities and Responsibility,' *Southern Journal of Philosophy* 31 (1993) 345-75; David Zimmerman, 'Acts, Omissions, and Semi-Compatibilism,' *Philosophical Studies* 73 (1994) 209-23; Alison McIntyre, 'Compatibilists Could have done Otherwise: Responsibility and Negative Agency,' *Philosophical Review* 103 (1994) 453-88; Randolph Clarke, 'Ability and Responsibility for Omissions,' *Philosophical Studies* 73 (1994) 195-208; Walter Glannon, 'Responsibility and the Principle of Possible Action,' *Journal of Philosophy* 92 (1995) 261-74; Alfred Mele, 'Soft Libertarianism and Frankfurt-Style Scenarios,' *Philosophical Topics* 24 (1996) 123-41; David Copp, 'Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility,' *Nous* 31 (1997) 441-58; Michael della Rocca, 'Frankfurt, Fischer, and Flickers,' *Nous* 32 (1998) 99-105; Michael Otsuka, 'Incompatibilism and the Avoidability of Blame,' *Ethics* 108 (1998) 685-701. For a bibliography of more articles, see John Fischer, 'Recent Work on Moral Responsibility,' *Ethics* 110 (1999) 93-139. See also the references in subsequent notes.

The Story of Jones and Black (and a brief history of a philosophical literature)

Frankfurt's argument against PAP is based on a simple thought experiment. It begins by inviting you to tell a story about an agent, Jones, who chooses to perform, and succeeds in performing, some action X. Tell the story so that it is vividly clear that Jones is morally responsible for doing X. If you are a libertarian, you may specify that Jones is an indeterministic agent who can choose otherwise, given the actual past and the laws. If you are a compatibilist, you may fill in the details so that Jones does X in a way that satisfies your favorite account of the counterfactual or dispositional facts that make it true that Jones could have done otherwise in the sense you think relevant to responsibility. Now, add to your story the following facts: there is standing in the wings another agent, Black. Black is interested in what Jones does. In particular, he wants Jones to do X and, moreover, Black has it in his power to prevent Jones from doing anything other than X.

Just how Black might force Jones to do X is, Frankfurt averred, not vital to the story. Perhaps Black is standing in the wings ready to offer Jones a coercive threat that would stampede Jones into doing X. Or maybe Black has a drug or a hypnotic procedure that would give Jones an irresistible desire to do X. Or maybe Black has a device in place which would directly affect Jones's nervous system in such a way that Jones's body would be forced to move, puppet-like, through an execution of X. Fill in the details as you like, so long as it is clear that Black can and would prevent Jones from doing anything but X.

The addition of Black to the story means that Jones could not have done other than X. But, Frankfurt argued, Jones is still responsible for doing X. After all, though Black could have intervened, he didn't. He didn't have to. Jones chose to do X and did X without any interference from Black. So the addition of Black to our story doesn't remove or in any way diminish Jones's responsibility for doing X.

Such is the recipe for telling a Frankfurt story.⁴ And such stories can be told. Stories, that is, in which everyone should agree that an agent is responsible for doing something even though everyone should also agree that the agent could not have avoided doing that thing. And such stories do indeed tell us something interesting about responsibility. Such stories show us that it is sometimes false that:

4 Or at least one kind of Frankfurt story. I will be arguing that this isn't the kind of Frankfurt story that we need.

PAP' A person is morally responsible for doing X only if that person could have done other than X.

Thus, whatever your conception of freedom, if you thought that John Wilkes Booth freely chose to kill Lincoln and hence is responsible, it would be absurd to change your mind if you discovered that there happened to be a Black-like figure waiting passively in the wings, prepared to force Booth's hand had Booth changed his mind. Given the choices he made and the actions he took, Booth is responsible for killing Lincoln even if he could not have avoided killing Lincoln.

This much everyone should agree to.⁵ We've always known that there is a gap between the choices an agent makes and the outcomes of those choices. Two drivers choose to drink and drive, but only one has a fatal accident; we hold the second driver responsible not just for his choice but also for the death regardless of whether he was able, at the time of the accident, to avoid causing death. And we've always held agents responsible for the outcomes of their actions even if those outcomes are overdetermined. Thus we all agree that Booth is responsible for Lincoln's death even if there was a back up assassin ready to kill Lincoln if Booth's attempt failed. Black is like the second assassin, except that he would operate by forcing Booth's hand should Booth fail to attempt the assassination. His presence overdetermines the action Booth chooses, but does not deprive him of responsibility for it. After all, in the actual world Booth acts as he does, not because of anything Black does, but because Booth *chooses* to act that way. And, we could add here, no matter what happens to come of that choice, Booth's choice is free and he is responsible for it precisely because there is at least one moment at which Booth *could have chosen otherwise*.

Reflection on cases of real life overdetermination illustrate how far **PAP'** is from **PAP**. In real life, we hold persons responsible for *outcomes* they could not have prevented, provided that these outcomes are the foreseeable consequences of actions they could have avoided. Frankfurt

5 Or at least everyone should agree that there is some natural way of individuating actions, events, and states of affairs such that a person may be responsible for an action, event, or state of affairs despite the fact that she could not have avoided performing that action or causing that event or state of affairs. It is, of course, possible to defend **PAP'** by insisting that in overdetermination or causal pre-emption scenarios events (actions, etc.) are individuated more finely and that the agent is responsible *only* for this more finely individuated event, action, or state of affairs. For an example of this way of responding to Frankfurt, see Carl Ginet, 'In Defense of the Principle of Alternate Possibilities: Why I Don't Find Frankfurt's Argument Convincing,' *Philosophical Perspectives* 10 (1996) 403-17.

stories show that we are also prepared to hold someone responsible for an *action* he could not have avoided, provided that the action is *the foreseeable consequence of a choice he could have avoided making*. This is enough to refute **PAP'**. But **PAP** says that an agent is responsible only if he could have done otherwise. Otherwise than he actually does. To refute **PAP** one would have to tell a story in which an agent, say Booth, is responsible for what he did even though he could not have done *anything* other than he actually did. But Frankfurt stories of the sort just told do not fit that bill. They do nothing to subvert the view that moral responsibility requires the agent's ability to do *something* otherwise — even if that is only to make a different choice.

What is needed to refute **PAP** and sustain Frankfurt's claim that alternatives are unnecessary for responsibility is a different kind of story. What's needed is a story in which Jones does X and is responsible for doing X, but we must concede that Jones cannot do *anything*, even deliberate, decide, or choose, other than he actually does.

Frankfurt himself never tells such a story. When speaking of **PAP** he speaks only of responsibility for overt actions, as opposed to mental acts like choices or decisions, and he seems to think that demonstrating the falsity of **PAP'** suffices for demonstrating the falsity of **PAP**.⁶ But he also seems to assume that such a story can be told. That is, he seems to assume that the schema for the thought experiment sketched above can be filled out in a way that ensures that Jones can neither do nor *decide* otherwise.

Whether or not that is so has turned out to be the central issue in the post-Frankfurt debate. More than thirty years later, no one has succeeded in telling a story that *uncontroversially* meets Frankfurt's specifications. Not that there has been a shortage of attempts. Lots of philosophers have been convinced by Frankfurt, and they have tried to convince the rest of us. Their stories typically depict Black as a highly skilled neurosurgeon who has cunningly inserted devices in Jones's brain that allow him to monitor and alter Jones's brain states without Jones ever noticing. Black can intervene, if need be, causing brain states in Jones which are the physical realizations of the choices and decisions Black wants him to make. But as luck would have it, Jones deliberates, decides, and acts in just the ways that Black wants, so Black never has to intervene. Since Black never intervenes, Jones's moral responsibility remains intact. But

6 For instance, he says: 'This, then, is why the principle of alternate possibilities is mistaken. It asserts that a person bears no moral responsibility — that is, he is to be excused — for having performed an action, if there were circumstances that made it impossible for him to avoid performing it' (Frankfurt, 8-9).

Black's intentions and power guarantee that Jones can neither do nor even decide otherwise.⁷

The philosophers who remain unpersuaded by these kinds of stories divide into two main camps. Some philosophers have argued that since the choices of an indeterministic free agent are not predictable in principle, Black cannot intervene until *after* Jones has chosen and this ensures that there is at least a moment during which an indeterministic Jones remains free to choose otherwise.⁸

Other philosophers do not appeal to the unpredictability of Jones's free choices but insist that even if Jones is predictable and controllable, Frankfurt stories, by their very nature, cannot rule out all of Jones's alternative possibilities. They argue that there is necessarily a difference between the causal history leading up to Jones's action in the actual scenario and the causal history leading up to Jones's action in the counterfactual scenario in which Black intervenes. Because there is this difference, there *is* something that remains up to Jones, even if it is only whether or not he does or chooses X with or without Black's intervention.⁹

There are replies to these sorts of arguments. Against the claim that indeterminism entails unpredictability, defenders of Frankfurt have argued that even the choices of an indeterministic agent might be reliably correlated with — even though not caused by — some prior 'blush' or other involuntary sign, and therefore predictable on that basis.¹⁰ Others have told stories in which Black has a godlike omniscience to predict what even an indeterministic agent will decide.¹¹

7 See, for instance, Daniel Dennett, *Elbow Room: The Varieties of Free Will Worth Wanting* (Cambridge, MA: Bradford Books 1984), 132; John Fischer, *The Metaphysics of Free Will* (Cambridge, MA: Blackwell Publishers 1994), chap. 7; Ishtiyaque Haji, *Moral Appraisability* (New York: Oxford University Press 1998), chap. 2.

8 Cf. David Widerker, 'Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities,' *Philosophical Review* 104 (1995) 247-61 and Robert Kane, *The Significance of Free Will* (New York: Oxford University Press 1998), 142-3 and 191-2.

9 See, for instance, Peter van Inwagen, *An Essay on Free Will* (Oxford: Clarendon Press 1983), chap. 5, and M. Naylor, 'Frankfurt on the Principle of Alternate Possibilities,' *Philosophical Studies* 46 (1984) 249-58.

10 See Fischer, *The Metaphysics of Free Will*, 136 and 145. See also Frankfurt, 6 n.3.

11 See, for instance, Alfred Mele, *Autonomous Agents* (New York: Oxford University Press 1995), 101, n.18 and 141-2, and Fischer, 'Libertarianism and Avoidability: A Reply to Widerker,' *Faith and Philosophy* 12 (1995) 119-25.

As for the second argument — that there is a difference between choosing X and choosing X because Black makes you, John Fischer has argued that this difference is not ‘robust’ enough to ‘ground’ attributions of responsibility. After all, if it is conceded that Jones’s choice *is* predictable because, say, Jones (deterministic or not) always blushes in a certain sort of way before he even begins to try to make his choice, then we can imagine a Black who can prevent Jones from even beginning to choose by watching for the blush. But in that case the *only* difference between what Jones actually does and what he would have done if Black had intervened is the blush or some other involuntary sign that Jones manifests before he even begins to try to decide what to do. Surely, Fischer argues, such a mere ‘flicker’ of freedom — the freedom to ‘do’ something that is not even an action — is not significant enough to be relevant to Jones’s responsibility.¹²

And so it goes. Libertarians insist, for the most part, that they are unconvinced by Frankfurt’s claim to have refuted P_{AP}. Compatibilists almost unanimously¹³ insist that P_{AP} is false and that Frankfurt has shown that the traditional debate is irrelevant. New and ever more arcane Frankfurt stories continue to be told. Inevitably, the discussion turns to an argument about which side has the burden of proof, always a sure sad sign of a philosophical impasse.

I think we should have avoided this mess. Things went wrong from the start. No one should ever have been persuaded by Frankfurt’s argument.

Before I can explain why, I need to draw a distinction between two very different ways in which Black might operate.

Two Ways of Getting Someone to Do What you Want

Suppose you want to ensure that someone does whatever you want him to do, but, like Black, prefer to avoid showing your hand unnecessarily. There are two different methods you might employ. For reasons that will become clear later on, I will call these ‘the method of conditional intervention’ and ‘the method of counterfactual intervention.’

12 Fischer, *The Metaphysics of Free Will*, 134–47

13 For one dissenting compatibilist, see Joseph Keim Campbell, ‘A Compatibilist Theory of Alternative Possibilities,’ *Philosophical Studies* 88 (1997) 319–30.

What makes someone a *conditional intervener* is the fact that his intervention is causally triggered by the beginnings of any action (overt or mental) contrary to the intervener's plan. If the subject begins to try or begins to do any undesired action, the intervener will prevent him from succeeding. The intervener might do this by hanging around, ready to leap into action at a moment's notice, or he might do it by setting out external and internal trip wires or barriers that will be activated if his subject starts to stray from the proper behavioral or psychological path. The key is that conditional intervention depends — depends causally — on the subject's beginning to try or beginning to do the wrong thing. To prevent his subject's misbehavior, the conditional intervener must set things up so that conditional devices are on the job *all the time*, ready to be triggered the moment the person begins to stray from the desired path.

What makes someone a *counterfactual intervener* is the fact that his intervention is causally triggered, *not* by the subject's trying or beginning to act contrary to the intervener's plan, but by some *earlier* event that is a reliable indicator of the fact that the subject will, in the absence of intervention, choose or act contrary to the intervener's wishes. This earlier event might be a blush, twitch, or other involuntary sign that occurs just before the subject begins to make an unwanted decision¹⁴ or it might be some event that takes place hours or even days earlier. In the limit case of a deterministic universe and a Laplacean predictor, it might be the state of the entire universe at some time before the subject was born.

What distinguishes these two methods is not the coercive machinery employed on those occasions when there is intervention. A conditional and a counterfactual intervener might both intervene by direct brain manipulation, for instance, or hypnosis, or the implantation of an irresistible inner compulsion, or any of the other methods that Frankfurt suggests. Nor does the distinction lie in the fact that the counterfactual intervener has an ability lacked by the conditional intervener: the ability

14 In my discussions of conditional and counterfactual intervention, I will be using 'decision' or 'choice' to refer to the mental acts that Black controls. But nothing hinges on this. If you have a different theory about the mental goings-on or acts that 'ground' or are relevant to an agent's moral responsibility, then feel free to make the relevant substitutions. Just remember that the difference between conditional and counterfactual intervention lies in the fact that the counterfactual intervener's interventions are causally triggered by something that occurs *before* the relevant mental act or goings-on even begins to take place, whereas the conditional intervener's interventions are triggered by the beginnings of the relevant mental act or goings-on.

to predict his subject's choices and actions. The conditional intervener might have such a predictive ability, but he does not rely on it in order to control his subject's behavior. What distinguishes these methods are two things. First, on those occasions when there is intervention, the counterfactual intervener intervenes earlier — before the subject even begins to make the wrong decision or perform the wrong action. Second, on those occasions when there is no intervention, different *causal counterfactuals* are true. In the case of the conditional intervener, it's true, for every action *X* or choice *Y* contrary to the intervener's plan, that if the subject had begun or tried to do *X* (choose *Y*), this would have triggered something which would have prevented him from succeeding. In the case of the counterfactual intervener, the relevant true counterfactual is a different one: for every action *X* or choice *Y* contrary to the intervener's plan, if the subject had shown some *earlier sign* that he was going to do *X* or to choose *Y*, this earlier sign would have triggered something which would have prevented him from even beginning to do *X* or to choose *Y*.

Note that these two methods are logically independent of each other. It's possible for someone to use the method of conditional intervention without also using the method of counterfactual intervention, and vice versa, and it's possible for someone to use both methods.

Finally, note that these methods are not *essentially* devices employed by one human being to control the choices and actions of another human being. I have described them this way because in the Frankfurt stories usually told in the literature Black is a human agent. But this is not vital. So long as the relevant causal counterfactuals are true, a machine or natural forces could occupy the role of either conditional or counterfactual intervener.

With this distinction in hand, we are in a position to take a more rigorous look at Frankfurt stories. Is Black a conditional intervener, a counterfactual intervener, or does he employ some combination of the two techniques? It's not always clear, since the literature has never recognized the distinction. In Frankfurt's original story, Black was explicitly a counterfactual intervener.¹⁵ On the other hand, in the first article published in response to Frankfurt's article, David Blumenfeld noted

15 [He] waits until Jones is about to make up his mind what to do, and he does nothing unless it is clear to him (Black is an excellent judge of such things) that Jones is going to decide to do something *other* than what he wants him to do. If it does become clear that Jones is going to decide to do something else, Black takes effective steps to ensure that Jones decides to do, and that he does do, what he wants him to do. Whatever Jones's initial preferences and inclinations, then, Black will have his way' (Frankfurt, 6).

that Frankfurt's way of telling the story seemed to beg the question against those who believe that the decisions of a morally responsible agent cannot be reliably predicted ahead of time. To remedy this defect, Blumenfeld proposed a story in which the role of Black is played by natural forces which are causally triggered by an event (the person's flushing deep red) which is caused by the person's decision not to perform a certain action. Blumenfeld's story, then, is one of conditional intervention.¹⁶ Subsequent Frankfurt stories told in the literature are of both varieties, and some are stories in which Black is both a conditional and a counterfactual intervener, ready to intervene both *before* and *after* Jones begins to decide or act contrary to Black's plan.

Understanding the distinction between conditional and counterfactual intervention is the key to understanding both the seductive charm of Frankfurt stories and why they ultimately fail.

In the first place, once we have the distinction in hand, it will quickly become apparent that, no matter how ingeniously or elaborately told, no story of conditional intervention can, *in principle*, provide a genuine counterexample to PAP.

Thus, tell a Frankfurt story of the sort common in the literature: Imagine Black as a nefarious neurosurgeon whose sensitive devices allow him to monitor every quantum twitch and quiver of every neuron in Jones's brain. Suppose, moreover, that Black has in place a host of remotely controlled nano-devices which can intervene at a moment's notice to change Jones's brain and mind the moment Black decides to intervene. As it turns out, Jones always chooses and acts just as Black would have him do, so Black never intervenes. But Black would have intervened if...

If what? How we answer this question determines whether our story is one of conditional or counterfactual intervention. If what would have prompted Black to intervene is some neural goings-on that would count as Jones's beginning to choose, or to try, or to do something contrary to Black's wishes, then Black is a conditional intervener. On the other hand, if Black's intervention would have been prompted by his prediction of how Jones would otherwise — without intervention — choose, and if Black's intervention would have operated before the unwanted choice had even begun, Black is a counterfactual intervener.

Now of course real life neuroscience cannot begin to parse such matters, but given that this is our fantasy we are allowed to stipulate whether Black's intervention would precede or follow the beginnings of

16 Blumenfeld, 'The Principle of Alternate Possibilities,' *Journal of Philosophy* 68 (1971) 339-45

choice. Some philosophers might insist that it makes no sense to speak of choices taking time, so that if a choice has begun it is already made. For them, the question shrinks to: Would Black have intervened before or after Jones makes his choice? Never mind. The question is whether Black's intervention would happen after the choice is made (or has begun to be made) or before. If Black is in a position to predict Jones's choices, and prepared to act, if necessary, before Jones even begins to go astray, he is a counterfactual intervener.

In the literature, Frankfurt stories about conditional interveners typically describe interveners who do their job by constant monitoring of their subject. But this is not necessary to a method of intervention's being conditional in the defined sense. What is necessary is that the person's attempts to do some contrary action would causally trigger something which would prevent the attempt from succeeding. So understood, it becomes apparent that stories of *conditional* intervention long antedate Frankfurt. The person who carried Locke's sleeping man¹⁷ to the locked room is a conditional intervener. If the man, on awakening, had tried to leave the room, this would have causally triggered something — the resistance of the locked door — which would have prevented his attempt from succeeding. But since he chooses to stay in the room of his own free will — and could have chosen otherwise — he is morally responsible for staying.

Locke's contented prisoner is already a counter-example, if one were needed, to **PAP'**. But because he is free to try to leave the prisoner is manifestly no counter-example to **PAP** itself. Locke's story can be made to sound more Frankfurtean by the addition of more apparatus. Imagine that the door is not locked, but that the jailer (Black) is standing by ready to lock it the moment the prisoner, Jones, begins to move towards the door. Or suppose that the latch will be tripped the moment that Black psychically detects the beginnings of Jones's resolution to leave the room. Or imagine, as a recent variation on a Frankfurt story has it, that Jones has placed locks on the neural pathways in Jones's brain in such a way that, while Jones is never *forced* to choose as he does, had his deliberations taken any other course, they — or rather their neurological realizations — would have found the alternative routes closed.¹⁸

17 John Locke, *An Essay Concerning Human Understanding*, bk. 2, chap. 21, sec.10

18 See, for instance, David Robb and Alfred Mele, 'Rescuing Frankfurt-Style Cases,' *The Philosophical Review* 107 (1998) 97-113, and David Hunt, 'Moral Responsibility and Unavoidable Action,' *Philosophical Studies*, forthcoming.

But, however we tell such stories, it should be clear that so long as Black's methods are *only* conditional, they *necessarily* fail to deprive Jones of all alternatives. Since Black's conditional devices are always on the job, it's true, at all relevant times, that if Jones began or tried to do anything contrary to Black's plan, he would be prevented from succeeding. But so long as the conditional devices are triggered *only* by the beginnings of Jones's contrary choices and actions, Black's mechanisms do not, and *cannot*, succeed in depriving Jones of his ability to choose or at least try or begin to choose otherwise. However we hedge him round, the conditional devices can only kick in if Jones makes a first wrong move. And that first move is — so far as the story can show — always up to Jones. One might well think that his capacity to make this first move, however quickly its upshots may be squelched,¹⁹ is precisely why Jones retains his moral responsibility.

Insofar as Black is a conditional intervener, he succeeds in depriving Jones of the alternatives required for *freedom of action*, but he *necessarily* fails to deprive Jones of the alternatives required for *freedom of will*. And, of course, it is freedom of will that is the classical locus of moral responsibility. To borrow Kant's words, Jones's 'will sparkles like a jewel in its own right' despite the fact that, due to a 'particularly unfortunate fate,' it is 'wholly lacking in power to accomplish its purpose.' The lesson of Frankfurt stories, insofar as they are stories about conditional interveners, is that the freedom required for moral responsibility is freedom of will, not freedom of action, and this freedom is not removed by the method of conditional intervention.

To show that PAP is false — that an agent may be responsible even when he cannot do *anything* other than what he does — Frankfurt needs to tell a story wherein Jones genuinely lacks freedom of will. He needs a mechanism which can rob Jones of the ability to even begin to choose or decide otherwise. Can a counterfactual intervener succeed where a conditional intervener cannot?

19 Note that it is only in the counterfactual scenario in which Jones's freedom of will is of brief duration. In the actual scenario — where Black does not intervene — Jones retains freedom of will at all relevant times. His freedom is fragile, insofar as he continues to enjoy it only so long as he continues to make the choices (form the intentions and resolutions, etc.) that Black wants him to make. But in this respect he differs from the rest of us in degree rather than in kind. For our freedom of will is also fragile; we continue to enjoy it only so long as we are not struck by a bullet, a fast car, or a stroke.

It has tacitly been assumed, by nearly all parties to the debate,²⁰ that counterfactual intervention works. That is, it has been assumed that if Black has the ability to predict Jones's choices and the ability to intervene ahead of time to keep things going his way, then Jones has *no* alternatives to doing exactly what Black wants with respect to the actions that he performs and the choices that he makes. It is for this reason, I suspect, that most incompatibilists think it so important to deny that Black *could* have this ability, and for this reason that most compatibilists think that they must concede that in Frankfurt stories Jones is unable to even begin or try to choose otherwise.

I will argue that this assumption is mistaken. I am not going to argue that the choices of a free agent are unpredictable. Nor will I claim that there is anything about choice or deliberation that a suitably empowered Black could not change. I will argue that *even if* Black has the *power* to predict and alter everything and anything that Jones thinks or does, this power does *nothing* to diminish Jones's ability to choose or act otherwise.

If that claim sounds remarkable, it shouldn't. Compare: The government has the *power* to keep me from doing all sorts of things by having me summarily thrown in jail. But I am nevertheless free to do what I like because the government has not *exercised* this power. What should sound remarkable is Frankfurt's claim that a counterfactual intervener, however powerful he is, manages to rob Jones of all alternatives without *ever* exercising his power. How can *that* be so?

It can't. To see why, let me tell you a story.

Heads I win, Tails you lose

In my story you and I make a bet on the outcome of the toss of a coin. I bet heads — I always bet heads; you bet tails. It comes up heads. I win. You pay up. The question is: Did I win fairly? Well, of course it depends on whether it was a fair coin and a fair toss. Let's stipulate that it was. That is, let's stipulate that there was nothing about the physics of the coin or its toss that made it more likely that the coin would come up heads

20 A possible exception is James Lamb ('Evaluative Compatibilism and the Principle of Alternate Possibilities,' *Journal of Philosophy* 90 [1993] 517-27). However, Lamb's criticism of Frankfurt's argument is marred by the fact that he fails to recognize the distinction between conditional and counterfactual intervention. In their response to Lamb, John Fischer and Paul Hoffman also fail to notice the distinction ('Alternative Possibilities: A Reply to Lamb,' *Journal of Philosophy* 91 [1994] 321-6).

rather than tails.²¹ The odds were 50/50 that this coin on this toss would come up heads, and in the course of the toss no outside forces intervened to change those odds.

In the ordinary course of things, that would settle the matter. You lost fair and square. But this is not an ordinary case; there is more to my story. I have a confederate named Black. Black is a mysterious figure, with unusual predictive powers. Black is able, somehow, to predict how any given coin toss will turn out if he does not intervene and his predictions are *always* right, no matter how fair the coin and its environment. How does he do it? Here people disagree. Some people think that Black is somehow sensitive to a sort of tickle or blush that, mysteriously, is a reliable indicator of how the coin will land (if he doesn't intervene).²² Others think that Black is just lucky.²³ Never mind. Let's stipulate that somehow Black has this power to predict the fall of the coin.

Thanks to his unusual predictive powers, Black also has the ability to act ahead of time in a way that ensures that the coin will come up the way he wants it to. Here's how he does it. Black makes his predictions in the morning; we never toss the coin until evening. This leaves Black with plenty of time to fix either the coin or its environment, if need be, in a way that nomologically guarantees that the coin will come up the

-
- 21 In saying this, I do not mean to imply that only a genuinely chancy coin toss is a fair one. But for the purposes of this story, I stipulate that this coin toss is fair in the most straightforward way; the laws of physics, together with facts about the coin and the environment in which it is tossed, assign an objective probability of .5 to the outcome heads and .5 to the outcome tails.
- 22 These people are in effect claiming, that, given the facts about Black's impressive track record, we are entitled to conclude that it's not just an accident that the coin always lands in the way that Black wants. They claim that the best explanation for the correlation between Black's desires and the outcome of the coin toss is that there is a lawlike but noncausal connection between some earlier event (the tickle or blush) and facts about what the outcome of any particular coin toss would be (in the absence of intervention). It is on the basis of this earlier event that Black makes his decision to either intervene or not.
- 23 These are people who argue that it's impossible for anyone, even someone with godlike knowledge, to *know* what the outcome of an indeterministic process would be (in the absence of intervention). However, even these people must I think concede that there is no contradiction in the following story: Black makes his predictions on the basis of a series of hunches about the fall of the coin and intervenes if and only if he predicts that the coin would otherwise land tails. The coin in fact always lands heads, about half the time due to objective chance (tested in the usual way), the other half due to Black's prior intervention (again, testable: he substitutes a weighted coin, or puts a magnet under the table). I thank Fred Dretske, Dagfin Follesdal, and others at Stanford for pressing this objection. See also notes 27, 28, and 29.

way that he wants. But Black also prefers not to show his hand unless he has to. If he predicts that the coin will come up the way he wants, then he does nothing. Either way, his job is done by noon. By the time we toss the coin, Black has retired for the day. You may imagine him far away, or fast asleep.²⁴

Finally, let's stipulate that Black is a friend of mine. He wants me to win. If he had predicted that without his intervention this particular coin toss would have come up tails, he would have fixed things so that it would have been nomologically necessary on this toss that the coin would come up heads.

But as a matter of fact, on that last toss, Black didn't intervene. He predicted (somehow) that the coin was going to come up heads, so he did nothing.

Now our question, once again, is: Did I win fairly or were you cheated? Knowing what you now know are you entitled to your money back? I don't think so. Here is my argument: Of course, I agree that *if* Black had intervened and forced the outcome of that toss, that would have been cheating. And, granted, Black was prepared to cheat on my behalf. But as it happens, I was lucky and Black didn't have to cheat. And because Black didn't intervene the coin toss was a genuinely chancy one, with .5 probability of coming up tails. I won fair and square.

Now if you don't buy this line of argument then it seems to me that you must think that Black's existence, in the background and at a distance, somehow affected the coin toss, changing it from one that could have gone either way to a toss that *had* to end as it did — heads. After all, you were prepared to agree, before you learned that Black existed, that the toss was fair, that it could have gone either way. But if you think that Black somehow makes a difference to the toss, then it seems to me that you are going to have to assert:

- (1) If Black had not existed, the coin might have landed tails.

I think (1) is false. Since Black did not *actually* cause the coin to come up heads, I think that the existence of Black is counterfactually irrelevant to how the coin lands. That is, I think that (1) is false and what's true is:

- (2) If Black had not existed, the coin would still have come up heads

I think this for the same reason that I think it's true that:

24 I stipulate this because I want to tell a story in which it's clear that Black is *only* a counterfactual intervener.

(3) If I had not bet heads, the coin would still have come up heads.

On the other hand, I agree that counterfactuals about indeterministic events are tricky. We evaluate counterfactuals by asking what happens at the possible worlds most similar to our own at which the antecedent is true. Theories of counterfactuals differ over what similarities are relevant.²⁵ You might argue like this: Given that this coin is genuinely indeterministic, there are worlds exactly like our world until just before the coin lands where it comes up tails. These are the similarities that count. Therefore, the truth of the matter is:

(1) If Black had not existed, the coin might have landed tails.

And of course if (1) is true, (2) is false.²⁶

I don't subscribe to this theory of counterfactuals myself, but I don't say that it is wrong. I don't have to, because this line of argument isn't going to help your case. If this is your reason for denying (2), then you are also going to have to say, for the same reasons that (1) is true, that it is also true that:

(4) If Clinton had not existed, the coin might have come up tails.

But, however good your theory of counterfactuals, blaming the president in this way won't get your money back in Vegas and it won't help here. Clinton didn't interfere with your coin; neither did Black.

On that last toss, you lost fair and square. Of course, knowing what you know now about Black and me, you would be crazy to keep playing this game. You would be crazy, because it is clear that so long as Black's powers remain intact, you *will* never win this game. The coin *will* never come up tails.

Does that mean the game is rigged? Well, yes and no. The complicated truth about this situation, it seems to me, is that on any given toss of the coin:

25 Cf. the debate between Michael McDermott ('Lewis on Causal Dependence,' *Australasian Journal of Philosophy* 73 [1995] 129-39, and 'Reply to Ramachandran,' *Australasian Journal of Philosophy* 74 [1996], 330) and Murali Ramachandran ('McDermott on Causation: A Counter-Example,' *Australasian Journal of Philosophy* 74 [1996] 328-9).

26 I am here assuming the standard reading of a 'might' counterfactual as equivalent to the negation of the corresponding 'would-not' counterfactual. (If p, it might be that q is true iff it's false that if p, it would not be the case that q.)

- (1) EITHER the coin comes up heads even though it could have come up tails OR the coin comes up heads and could not have come up tails.

It follows from this that the coin *will* never come up tails, not that it *can't*. On some tosses the coin can up tails as easily as it can come up heads. But, thanks to the peculiar setup, it so happens that the coin *can* come up tails only on those occasions that it actually comes up heads.

When Black interferes, the game is rigged. The chance that the coin will come up heads on *those* occasions is 1. But on those other occasions when Black does not intervene, the objective probability of tails remains .5. Unluckily (for you) those are just the occasions on which the coin happens to come up heads. If we continue to bet, you will continue to lose, but you will have been cheated, probably, only about half the time.

When I've tried to explain these complicated facts to people I have sometimes heard the following objection:

Your story makes no sense. You say that the outcome of the coin toss is genuinely chancy; this means that given the laws and the initial conditions, the coin could have landed either way. But you also claim that Black knew ahead of time that it would land heads.

If my objector is puzzled about how Black can be such a reliable predictor of an event which lacks deterministic causal antecedents, I am sympathetic. But his objection is to Frankfurt, not to me. If it's impossible for Black to know the outcome of an indeterministic coin toss, then it's also impossible for Black to know the outcome of an indeterministic process leading to decision or choice.²⁷ And if that's so, then counterfactual intervention is not possible after all, and Frankfurt has failed to undermine the traditional debate.

Because I think that Frankfurt's argument goes wrong in a more fundamental way,²⁸ I prefer not to rest my objection on the impossibility of Black knowing the outcome of the coin toss.²⁹

27 Some supporters of Frankfurt appeal to the medieval doctrine of 'middle knowledge' which was invoked by Molina to explain God's knowledge of so-called 'counterfactuals of freedom.' Molina claimed that there are true counterfactuals about how an indeterministic free agent would choose in nonactual circumstances and that God's knowledge includes knowledge of these counterfactuals. See John Fischer, 'Libertarianism and Avoidability: A Reply to Widerker,' *Faith and Philosophy* 12 (1995) 119-25. For a critical discussion of this view, see William Hasker, *God, Time, and Knowledge* (Ithaca, NY: Cornell University Press 1989), 15-52.

28 That is, I think that Frankfurt's argument fails regardless of whether Jones's choice

Another objection I have heard goes like this:

The coin can't come up tails on *any* toss. On any given toss, if the coin is going to come up tails, Black will make it come up heads.

To this I say: Your argument contradicts itself. If it is *literally* true that the coin is going to come up tails, then it will come up tails, no matter what Black does or doesn't do.

This usually provokes something like the following:

That's not what I meant. I meant that if the coin is going to come up tails *unless Black intervenes*, then Black will make it come up heads. So one way or another, it's always true that the coin will come up heads.

To which the reply is: I agree that the coin *will* always come up heads. And I agree that whenever Black intervenes the coin cannot come up

is an indeterministic event or a deterministic event (which may unproblematically be known in advance). The coin story could be reformulated as the story of a deterministic but genuinely fair coin toss. Roughly, a deterministic toss is a fair one if the coin can land either way given the laws and intrinsic facts about the coin, its environment, and the abilities of the tosser immediately before the coin toss. The last part is meant to rule out cases where the tosser has the ability to reliably bring it about that the coin lands a certain way. Even though, as Black's accomplice, I know THAT the coin will land heads every evening at the appointed time, it's false that I know HOW to make the coin land heads. If you asked me to do it again, I would be as likely to fail as succeed.

- 29 Do I need to stipulate that Black has knowledge, or does it suffice to say that somehow the coin *always ends up landing* in the way that Black wants it to land, whether because of Black's knowledge or because Black is an uncommonly lucky guy; a guesser with a 100% success rate?

I'm not sure, because my aim is to construct a story with the same formal structure as a Frankfurt (counterfactual intervention) story and it's not clear what this is. Does the intuitive force of Frankfurt stories depend on Black's *knowing* what Jones will decide? Or is it enough that somehow Jones *always ends up deciding* in the way that Black wants him to decide, whether because of Black's knowledge or because Black is an extraordinarily reliable guesser? Or could it be that what makes the stories work is that we equivocate between these two ways of thinking of Black's powers? When we think of Jones as a morally responsible agent, we think that Black is just a *very* lucky guesser; when we think of Jones as someone who lacks alternatives, we think that Black *knows* his future choices.

I leave this as a question to be debated among the followers of Frankfurt. In my criticism of Frankfurt I assume what Frankfurt assumed and what makes his case strongest; that Black knows what the outcome of an indeterministic process would be (in the absence of intervention). But my criticism stands intact if Black's success is due to an incredible run of luck.

tails. But this doesn't prove what you are trying to prove: that on each and every coin toss, the coin *cannot come up tails*.

Another objection is more subtle: 'Granted, the coin was in fact not weighted nor was the environment in fact rigged in any other heads-determining way. But these facts do not suffice for it's being true that the coin (on this toss) could have landed tails. The coin could have landed tails only if the following *counterfactual* is true: 'If the coin were about to land tails, no outside force would make it land heads.' And this counterfactual is false. For if the coin were about to land tails, Black would have predicted this and intervened. And if Black had predicted this and intervened, the coin would be forced to land heads. So if the coin were about to land tails, it would be forced to land heads.'³⁰

But this objection relies on a form of counterfactual reasoning which is generally considered invalid: hypothetical syllogism. An example: 'If I jumped off this bridge, I would have arranged to be wearing a parachute. If I were wearing a parachute, I would not be killed. So if I jumped off this bridge, I would not be killed.'³¹

Another argument I have heard goes something like this:

Given the facts of the story the coin *cannot* come up tails. Proof. Assume for the sake of reductio that:

(1) The coin comes up tails.

Now in the story as told

(2) If the coin would come up tails without Black's intervention, then Black would intervene and ensure that the coin does not come up tails.

(3) Ergo: the coin does not come up tails.

Which is a contradiction, hence (1) cannot be true.

30 Cf. Fischer: 'If he were about to refrain (in the absence of intervention by an external agent or factor) the triggering event would already have occurred. [If the triggering event had already occurred, Black would have intervened and forced Jones to act, in which case Jones would not have been able to refrain.] If Jones were about to refrain, he would be rendered unable to refrain' ('Alternative Possibilities: A Reply to Lamb,' 326).

31 P.J. Downing was the first to direct our attention to an instance of this fallacy in his 'Subjunctive Conditionals, Time Order, and Causation,' *Proceedings of the Aristotelian Society* 59 (1958-59) 125-40. For discussion of this and other counterfactual fallacies, see David Lewis, *Counterfactuals* (Cambridge, MA: Harvard University Press 1973), 31-6. See also Jonathan Bennett, 'Counterfactuals and Temporal Direction,' *Philosophical Review* 93 (1984) 57-91.

The peculiar appeal of this argument illustrates the dangers of using *reductios* in subjunctive arguments. When we give a *reductio* argument we are attempting to make a counterfactual claim. We are saying that if a false proposition were true, then CONTRADICTION. It's easy to forget that in this counterfactual situation, some counterfactuals that are actually true may be false. And that's what's gone wrong here. This argument contradicts itself, but not in the way *reductios* are supposed to. The argument from (1) and (2) does indeed entail (3). But that is because (1) and (2) are inconsistent. There is no possible world at which Black has the knowledge, intentions, and powers that (2) ascribes to him and at which the coin comes up tails. That is, the only worlds at which the coin comes up tails are ones at which (2) is false and Black either makes a false prediction, or lacks the intention to make the coin land heads, or his coin-controlling powers fail. Doesn't that show that our original story was already inconsistent? No, for in our story the coin will never come up tails.

This reply sometimes provokes something like the following: 'You are equivocating on senses of 'can.' You are pointing out that it is *logically possible* for the coin to come up tails. But what's issue is whether it's possible in some stronger sense — whether it's possible given all the laws and all the relevant facts about the coin, including the facts about Black's knowledge, intentions, and powers.'

I deny that I'm equivocating. I insist that the coin can come up heads in exactly the sense that I specified when I first told you the story — the coin toss is, by stipulation, a genuinely chancy event. It has a .5 objective probability of coming up tails. Furthermore, it is also true that the coin satisfies our ordinary (looser) standards of being a coin that could have come up tails. Given the laws and given all the relevant *intrinsic* facts about the coin and its environment just before it is tossed, the coin could have come up tails.

At this point there is, so far as I can tell, only one line of argument open to you. It goes something like this:

Look. You must agree that, given the facts of the story, we know that one way or another the coin will come up heads. Given that this is so, given that this outcome is a *foregone conclusion*, we must say that the coin could not have come up tails.

To see the problem with this argument, consider that whenever *any* coin is tossed, however fairly, and it comes up heads, it is then *literally* a foregone conclusion that it comes up heads. Should we conclude that no coin that comes up heads could have come up tails? Some see no fallacy in this way of arguing. These people are called fatalists. They think that nothing could ever happen otherwise than it does. They think that it follows, from the mere fact that this evening's coin *will* land heads, that

it *must* land heads.³² If you are a fatalist, then nothing I can say is likely to change your mind. But I'm not going to give you your money back.

Conclusion

The moral of my story is, I hope, by now apparent. If you are a libertarian, it will be an especially easy moral to draw. Simply substitute 'Jones chooses to do X' for 'the coin comes up heads' and everything just now said carries through without revision.³³

For the compatibilist, the story is complicated only by the need to provide a compatibilist analysis of 'could have chosen otherwise.'³⁴

32 I take a fatalist to be someone who argues, on the basis of considerations of truth and logic alone, that everything that happens, including everything we do, is necessary (unavoidable, could not have been otherwise, etc.). A classic form of fatalist reasoning is as follows: 'It's true that I will eat corn flakes for breakfast tomorrow. Necessarily, if it's true that I will eat corn flakes for breakfast tomorrow, then I will eat corn flakes for breakfast tomorrow. So I *must* eat corn flakes for breakfast tomorrow.' A slightly more subtle form of fatalist reasoning is represented by the following argument: 'I can eat pancakes for breakfast tomorrow only if I can do so, given all the facts. But all the facts include facts about the future, including the fact that I will eat nothing but corn flakes for breakfast tomorrow. So I cannot eat pancakes for breakfast tomorrow.' For a good discussion of other fatalist fallacies, see David Lewis, 'The Paradoxes of Time Travel,' *American Philosophical Quarterly* 13 (1976) 145-52, reprinted in his *Philosophical Papers Volume 2* (New York: Oxford University Press 1986), 67-80.

33 Of course a libertarian need not suppose that every choice is 50/50, and there can be room for the agent's past to weight the odds, but it should be clear that the preceding argument carries through however the odds are structured.

34 I don't mean to suggest that this is easy. Just before Frankfurt published his article, there was an enormous literature devoted to the question of whether any so-called 'conditional analysis' of 'could have done otherwise' could be defended against various sorts of counterexamples, and the prospects for a successful account were beginning to look bleak. But that was thirty years ago. We now have a few advantages we didn't have then, most notably, the Lewis-Stalnaker semantics for counterfactuals. See Lewis, *Counterfactuals*, and also 'Counterfactual Dependence and Time's Arrow,' *Philosophical Papers*, vol. 2, *ibid.*, 3-66 and Robert Stalnaker, *Inquiry* (Cambridge, MA: Bradford Books 1984), esp. chaps. 6-8. There is also the question of the relation between abilities, capacities, dispositions, and counterfactuals. The old literature didn't draw any distinctions; it was assumed that to say that someone/something has an ability, capacity, or disposition to (do) x is just to say that she/it would (do) x, given certain conditions. This assumption is no longer accepted for dispositions like fragility and solubility, and there is an interesting literature of puzzle cases which are similar, in many ways, to Frankfurt stories. I suspect that this literature has something to tell us about how we should think about the abilities and capacities that constitute free will. See C.B. Martin, 'Dispositions

Given any plausible analysis of this modal locution, the moral is the same as it is for the libertarian. The presence of a counterfactual intervener of the sort Frankfurt described does not rob a free agent of the ability to choose otherwise. The complicated truth about Frankfurt's Jones is that whenever Jones chooses X:

- (2) EITHER Jones chooses X but could have chosen not-X OR Jones chooses X and could not have chosen not-X.

So long as Black does not intervene, Jones can choose to do what he likes. It so happens, thanks to the peculiar setup of the case, that Jones can choose to do otherwise only if he chooses what Black wants him to choose. So long as Black does not interfere, Jones can choose otherwise, even though he doesn't. He has alternatives even though he does not take them.

More austere, the point is this. The inference from:

(P and Possibly not-P) or (P and Necessarily P)
to
Necessarily P

is fallacious. To suppose otherwise is to permit the inference from:

P
to
Necessarily P

And this is the logic of the fatalist.³⁵

Received: July, 1998

Revised: February, 1999

Revised: December, 1999

and Conditionals,' *Philosophical Quarterly* 44 (1994) 1-8, and David Lewis, 'Finkish Dispositions,' *Philosophical Quarterly* 47 (1997) 143-58. I thank David Sanford for drawing my attention to this literature.

- 35 Earlier versions of this paper were presented at the University of Southern California, Stanford University, and at the 1999 Pacific Division meetings of the APA. I thank audience members for their questions and helpful criticism. I'm also grateful to Mark Bernstein, Randolph Clarke, Ishtiyaque Haji, Frances Howard-Snyder, Hud Hudson, Robert Kane, Tomis Kapitan, Ned Markosian, Alfred Mele, Adam Morton, Murali Ramachandran, David Sanford, Howard Sobel, Dan Speak, Tim O'Connor, Peter van Inwagen, David Widerker, and two anonymous referees for the *Canadian Journal of Philosophy* for written comments and helpful discussion. My deepest thanks and appreciation are to Terrance Tomkow.