

John Martin Fischer,
The Metaphysics of Free Will
(Oxford: Blackwell, 1994)

KADRI VIHVELIN
University of Southern California

The two most important arguments concerning freedom, moral responsibility, and determinism during the last quarter century are the so-called “modal argument”, articulated independently by Ginet and van Inwagen for the thesis that the freedom to do otherwise is incompatible with determinism, and Frankfurt’s famous argument for the thesis that we may be morally responsible even if we are unable to do otherwise.¹ Fischer has been convinced by both arguments. He calls himself a “semi-compatibilist”—“compatibilist” because he holds the traditional compatibilist view that causal determinism is compatible with moral responsibility; “semi” because he agrees with the incompatibilist thesis that determinism rules out *all* alternative courses of action, entailing that we can *never* do otherwise. Despite this concession to the incompatibilist, Fischer’s bottom line is that we’ve got the kind of freedom we most care about—roughly, the kind that’s a necessary condition of rational choice, personhood, and moral responsibility—even if the modal implications of determinism are every bit as bad as the incompatibilist says they are.

Fischer’s book is well worth reading for a number of reasons, but perhaps its most important virtue is that it brings together two sets of issues that tend not to be discussed together, and tend not to be discussed by the same people. Frankfurt’s counterexamples to the thesis now widely known as the Principle of Alternate Possibilities (the thesis that someone is morally responsible for what he’s done only if he could have done otherwise) have achieved almost the status of Gettier counterexamples in the free will literature.² The intuitions invoked by Frankfurt stories are highly compelling, and almost everyone agrees that they show us something important about moral responsibility, even though it’s still hotly disputed what this is. The modal argument, by contrast, has been influential

chiefly among incompatibilists. Why the disparity? I suspect the answer is that Frankfurt got there first. His article was written in 1969, and by the time the modal argument appeared in the late seventies, many philosophers had been persuaded that Frankfurt had shown that, at least so far as moral responsibility for *acts*³ is concerned, determinism poses no threat. But even if this is right, it remains a separate and interesting question whether there are good reasons for believing that determinism entails that we are never able to do anything other than what we actually do. John Fischer is one of the few (semi) compatibilists to take seriously this metaphysical and modal question.

There is more to the book than just a defence of incompatibilism about freedom to do otherwise and compatibilism about moral responsibility. There is a discussion of the problem of free will and divine foreknowledge, a discussion of Newcomb’s problem, and a discussion of different theories of counterfactuals and their relevance to the various free will problems as well as to Newcomb’s problem. Fischer’s view about divine foreknowledge is parallel to his view about determinism: he thinks that divine foreknowledge is incompatible with the freedom to do otherwise but compatible with moral responsibility. His position on Newcomb’s problem is that it makes a difference whether the predictor is merely inerrant or infallible. In the case of a merely inerrant predictor, Fischer is a two-boxer; in the case of an infallible predictor, Fischer is a one-boxer. Space does not permit me to discuss Fischer’s treatment of these two problems. I will restrict my comments to the issues which take up the substance of the book—Fischer’s defence of incompatibilism about free will and his defence of compatibilism about moral responsibility.

I will begin with Fischer’s discussion of Frankfurt’s well-known counterexamples to the Principle of Alternate Possibilities (PAP). The recipe for constructing Frankfurt counterexamples to PAP goes as follows. Imagine someone—call him Jones—who deliberates, decides, and does something *x* and who also satisfies whatever further conditions you think are required to make it true that he is morally responsible for doing *x*. (If you are a libertarian, you may stipulate that Jones’ decision is not causally determined, that Jones can choose otherwise, given the actual past and the laws, that Jones makes things happen by way of agent-causation, and so on.) Then add to the story the following further facts: There is someone—call him Black—who’s got the power and the intention to make Jones do *whatever* he wants him to do, but who prefers not to show his hand unnecessarily. By happy co-incident, Jones decides to do exactly what Black wants, so Black does not intervene and Jones does *x* on his own without ever becoming aware of Black’s existence. Jones is still responsible for doing *x*. But Black’s power and intentions ensure that Jones could not have avoided doing *x*. So PAP is false.

Up to this point, just about everyone agrees.⁴ But what’s the relevance of this for the debate concerning moral responsibility and determinism? Fischer argues that the only plausible reason for thinking that determinism is incompatible with moral responsibility is that deterministic causal conditions deprive us of alterna-

tive courses of action, rendering us unable to do other than what we actually do. He argues that Frankfurt stories show, not just that being able to do otherwise is not necessary for responsibility, but also that the facts that suffice for moral responsibility are facts about what happens in “the actual sequence”—facts about the agent’s actual reasoning, decision, and action. Perhaps it’s often, typically, or even always true that someone who decides, for reasons of her own, to do something is also able to do otherwise (as libertarians and traditional compatibilists believe); perhaps we merely falsely think that this is true (as hard determinists believe). But the case of Jones shows that this “extra fact” is irrelevant to questions about moral responsibility. Jones is responsible, even though he could not have done otherwise, because his action was caused by his own deliberation and choice, in the right kind of way. What counts as the “right kind of way”? Fischer offers us the following sketch (and promises a more detailed account in a forthcoming book): An action is caused in the right kind of way just in case it’s caused by a “weakly reasons-responsive mechanism”; that is, just in case there is a nomologically possible world (eg. a world without Black in the wings; eg. a world with a different past) in which the same mechanism operates, there is sufficient reason to do otherwise, and the agent recognizes the reason, chooses, and acts on it. (166)⁵

Fischer’s way of defending (semi) compatibilism is appealing. But I think it’s a mistake to think that Frankfurt-style stories can undercut the traditional debate about determinism and moral responsibility. To see why, let’s remind ourselves what the traditional debate was about. The traditional debate was about whether deterministic causal conditions entail that we can never do *anything* other than what we actually do; the assumption shared by both sides was that if this is so, then we are never morally responsible for anything. The standard compatibilist opening move, in the pre-Frankfurt literature, was to argue that when we say that someone could have done something (in contexts where questions of responsibility are at stake), what we mean is that she would have done it if she had chosen or tried to do it. The standard incompatibilist reply was to point out that the truth of this conditional provides, at best, only a *necessary* condition for the truth of ‘she could have done it’; someone who can’t do something because she’s unconscious (or under hypnosis or suffering from a phobia or other psychological impediment) may nevertheless be such that *if* she chose or tried to do it, she would. Both sides agreed that these kinds of cases show that someone may be unable to do something *because* she’s unable to choose or try to do it.⁶ Incompatibilists then argued that causal determinism entails that no one is ever able to choose or try to do anything other than what they actually do. And compatibilists were left on the defensive, searching for an account of freedom of will which would draw the intuitively correct distinctions between those unable to will or choose otherwise and the rest of us.

Frankfurt’s aim was to restore the advantage to compatibilists, at least so far as moral responsibility is concerned, by showing that the traditional dispute about whether determinism renders us unable to do otherwise is irrelevant to the ques-

tion of responsibility. It’s irrelevant, according to Frankfurt, because it rests on the false assumption that someone is responsible for what she did only if she could have done otherwise. But, as we’ve just seen, the real issue is whether determinism renders us unable to do otherwise by rendering us unable to *choose* otherwise. So Frankfurt stories must do more than convince us that Jones is morally responsible even though he cannot do otherwise. We must also be convinced that Jones is responsible even though he cannot *choose* otherwise.

You might be tempted to think that we can tell a Frankfurt story that succeeds in convincing us of this. For can’t we just stipulate that Black’s power over Jones extends, not just to the movements of Jones’ body, but also to everything that happens inside Jones’ head? If we tell such a story, then it seems that Jones cannot avoid carrying out Black’s plan for him, down to the smallest detail, including how he deliberates and what he decides, chooses, and tries to do. But since we can also stipulate that Black never lays a finger on Jones, it seems that Jones is still responsible for what he does.

But even though we can tell Frankfurt stories in which Jones will *end up* choosing and doing exactly what Black wants him to do, we cannot tell Frankfurt stories in which Jones is unable to at least *try* to deliberate, decide, choose, and do otherwise. Stories which succeed as counterexamples to PAP (or to related principles about choice, decision, deliberation, etc.) are stories in which Black remains in the wings, playing no causal role in anything that Jones does (thinks, decides, etc.), and in which Black’s remaining in the wings causally and counterfactually depends on what Jones does (thinks, decides, etc). But these facts also guarantee that Jones remains able to at least *try* to do (deliberate, decide, choose, etc). otherwise.

To see why, consider this way of spelling out the details of a Frankfurt story. (This is the way Fischer in fact tells the story.) Black’s power over Jones takes the form of a device that Black has installed in Jones’s brain, a device Black can use to monitor and, if necessary, to directly manipulate Jones’ brain states. At the actual world, Jones deliberates and chooses exactly as Black wants him to, so Black never has to intervene. But if the information from Jones’ brain had been different, then Black would have pressed the button, activating the brain manipulation machinery, causing Jones to once again deliberate, decide, etc. exactly as Black wants him to. But because causal signals take time, there is a time interval between the occurrence of each event that would trigger Black’s intervention and the first brain event that would (in the counterfactual scenario) be caused by Black’s manipulation. This means that Frankfurt stories include a sequence of ineliminable moments of freedom—moments during which Jones could have done something differently—deliberated, decided, chosen and/or acted differently (or at least *tried* to deliberate, decide, choose or act differently).

Fischer suggests that Black may have the power to prevent Jones from even trying to deliberate or decide differently. Suppose that Jones is so highly predictable that Black knows, before Jones starts to deliberate, what the upshot of deliberation will be. Then it seems that Black’s power includes the power to stop

Jones from even beginning to deliberate in a way contrary to what Black wants. If so, then while it may remain true that something different could have *happened* to Jones, it seems that Jones could not have *done* anything different. (136)

But there is a counterfactual fallacy lurking here. Let's grant, for the sake of argument, that Black's power includes the power, at t_1 , to correctly predict (on the basis of a blush or some other involuntary body movement) how Jones will deliberate and decide during the time interval t_3 – t_5 .⁷ Then it's true that, *if* Black had predicted, at t_1 , that Jones would (in the absence of interference) deliberate and decide "against Black's will" during t_3 – t_5 , then Black *would* have pressed the button at t_2 , forcing Jones to deliberate and decide the way the way Black wants. That is, Black's predictive powers, together with the brain manipulation machinery, entail that Black *could* have made it true that Jones lacks even a moment of freedom. But it's a mistake to infer "S made it the case that P" from "S could have made it the case that P". What *actually* happened was that Black predicted that Jones would freely deliberate and decide in the way Black wanted him to. So Black never actually pressed the button, and Jones remained able, during the time interval t_3 – t_5 , to at least begin to deliberate contrary to Black's will.⁸

The moral is this: If we begin by thinking of Jones as someone who has freedom of will, then we cannot, contra Frankfurt, remove this freedom by adding Black to the stories. So long as Black *remains in the wings*, his power over Jones is necessarily limited. Black has the power to prevent any attempt by Jones to act (deliberate, decide, etc.) against Black's will from being successful, but Black lacks the power to prevent Jones from making these attempts in the first place. Jones' will, though wholly lacking in power to achieve Jones' ends (except insofar as they co-incide with Black's) remains as free as it was before the introduction of Black, and continues "to sparkle like a jewel in its own right" (to use Kant's famous words).⁹

Although Fischer might quarrel with my diagnosis of why Black cannot deprive Jones of all his freedom, he agrees that Frankfurt stories include what he calls "a flicker of freedom". (145) And he agrees that if the incompatibilist is right, then causal determinism extinguishes "not just a prairie fire of freedom, but also the tiniest flicker". (135) But he nevertheless argues that Jones's flicker of freedom is irrelevant because it's "not sufficiently robust" to "ground" attributions of moral responsibility. (140)

This is puzzling. No one ever claimed that having alternatives "grounds" responsibility in the sense of providing a condition that is both necessary and sufficient for responsibility; the claim was only that having alternatives is a *necessary* condition of moral responsibility. To show that this claim is false, we need to tell a story about someone who is responsible even though she has *no alternatives*—not even momentary "flickering" alternatives.

Still, Fischer seems to be on to something when he complains that Jones' mere "flicker" of freedom is not relevant to his responsibility for his action. Here's a way of spelling out what I think his worry really is. Compare Jones to two other

agents, Abel and Unlucky. Abel is the person we thought of in the first stage of the Frankfurt thought experiment, before we introduced Black. He is the paradigm case of someone with freedom of choice and action, who acts freely and is responsible for what he does. Jones and Unlucky are Abel's psychological twins, with the same values, standing beliefs, character traits, and so on. Let's pick a particular occasion on which all three have the same beliefs and desires, deliberate in the same kind of way, and decide, for their own reasons, to do something x instead of something else y . (To make it vivid, suppose that "x" is staying at home and working, and "y" is driving ten miles to the university to meet a student.) But here the similarities end. Black doesn't care what Abel does, but Black has a plan for both Jones and Unlucky. By happy co-incidence, Jones decided to do what Black wanted him to do, so Black does not interfere and Jones, like Abel, succeeds in acting on his decision and is morally responsible for doing x . But Unlucky's decision did not co-incide with Black's plan, so Black presses the button activating the brain manipulation machinery and Unlucky ends up "changing his mind" and doing y instead.

Unlucky is not morally responsible for doing y . But note that there is no difference between Jones and Unlucky so far as alternatives go. Both Jones and Unlucky lack any significant alternatives; neither can avoid doing what Black wants them to do. And both Jones and Unlucky retain a flicker of freedom; the time at which Unlucky decided to do x was a time at which it was true that he could have decided (or tried to decide) otherwise. This suggests that something other than alternatives is doing the work in explaining why Jones is responsible and Unlucky is not responsible. And it seems to support Fischer's claim that what's relevant, so far as responsibility is concerned, is what happens in "the actual sequence". Intuitively, the difference between Jones and Unlucky is that Jones's action (staying at home and working) was caused in the way that Abel's action was caused—the *right* kind of way, by a "weakly reasons-responsive mechanism" (practical reason), whereas Unlucky's action (driving to the university to meet the student) was caused in the *wrong* kind of way, by a mechanism which isn't weakly reasons-responsive (Black's direct manipulation of Unlucky's brain states).

But this is going too fast. I agree that Unlucky isn't morally responsible for his action, and I agree that we think he isn't morally responsible because his action wasn't caused "in the right way". But I don't think that the case of Unlucky supports Fischer's semi-compatibilism. On the contrary, closer examination will reveal that it supports our traditional pre-Frankfurt intuitions about moral responsibility. I leave it an open question whether these intuitions are ultimately defensible and whether they can be reconciled with a belief in determinism. My point is only that there are no shortcuts to be had in this area; we cannot avoid the traditional debate about determinism and freedom to do otherwise.

First point: Note that Unlucky remains responsible for *something*—he is responsible for what Kant would call his "pure" act of will. Unlucky decided to stay home and work, then Black pressed the button and Unlucky ended up driving in

choice nor her action would have caused this to be the case. To put it another way, if determinism is true, then we have abilities that we would exercise only if either the past or the laws were different.

Fischer is no longer confident of this reply. The issue, as he now sees it, is whether a *merely counterfactual* power over the past or the laws gives us grounds for rejecting compatibilism. That is, he thinks that there are intuitive grounds for inferring from "if S were to do x, the laws (past) would be different" to "S cannot do x" *even if* we are clear about the fact that we are not attributing any incredible causal powers to S. (67, 78)

Here's a way of motivating his worry. Consider a prisoner locked in a room. If she left, the door would have to be unlocked or somehow blasted open—there's no other way for her to leave. But the prisoner lacks the power to cause either to be the case. Therefore, the prisoner cannot leave the room. The incompatibilist reasons in a similar way about the past and the laws. If a deterministic agent did x (for any action x she doesn't in fact do), then either the past or the laws would have to be different—there's no other way for her to do x. But no human being has the power to cause either to be the case. Therefore, the deterministic agent cannot do anything she doesn't in fact do.

There are, of course, obvious differences between the prisoner and the paradigm example of a rational and apparently free deterministic agent. And one might think that a good compatibilist strategy would be to say what these differences are, and to argue that *these* differences, as opposed to the mere fact of deterministic causation, are what's relevant to the question of what we are able to do. Somewhat surprisingly, Fischer doesn't consider this way of defending compatibilism. His discussion remains at a more abstract level, focusing on the relevance of rival theories of counterfactuals to the free will/determinism question. In particular, he gives an excellent exposition of Lewis's theory of counterfactuals and of Lewis's use of his theory as part of a defence of compatibilism in response to the Conditional argument. Lewis argues that a deterministic agent has the following ability: the ability to do something x (where x is something the agent doesn't in fact do) such that if she did x, then *most* of the past prior to her action would have been *exactly* the same, and the laws would have been *nearly* the same, but there would have been a tiny "divergence miracle" shortly before she does x.¹² Fischer doesn't argue that Lewis's way of defending compatibilism is unviable or doomed to failure. Nor does he argue that a compatibilist reply based on a rival theory of counterfactuals (a theory which tells us to evaluate counterfactuals by considering *only* those worlds with the same laws as our world) is untenable. Nevertheless, he concludes his discussion by declaring that we have reached a "dialectical stalemate" due to hopelessly conflicting intuitions about the relation between counterfactuals and 'can' claims. The stalemate is this: The incompatibilist thinks that the truth of "If S did x, then the laws (past) would be different" suffices for the truth of "S cannot do x", whereas the compatibilist thinks that the truth of these counterfactuals is consistent with the truth of "S can do x."

This doesn't seem right. If there is a compatibilist account of 'can do x' which gets the *uncontroversial* cases right (eg. someone in a coma cannot raise her hand), then we have grounds for rejecting the inference from "if S did x, the laws (past) would be different" to "S cannot do x". And in the absence of any *further* reasons for thinking that deterministic causal laws are, in the relevant respects, like chains or prison bars, the stalemate has been resolved in favor of the compatibilist.¹³ Fischer doesn't offer any general argument against the possibility of a compatibilist account of "can do x"; on the contrary, he provides a sketch of what such an account might look like. In an earlier chapter he argues against a view he dubs "Restrictivism" (defended by van Inwagen, Susan Wolf, and Dennett, among others)¹⁴, which says that *even if* determinism is false, there are many occasions on which we cannot do otherwise (since we can't do what we think is crazy, morally indefensible, and so on.) Fischer argues against Restrictivism by defending the following account of "S can do x": 1) S can think of reasons that will generate a desire to do x; 2) S can try to act on this desire; and 3) If S were to try to act on this desire, S would succeed. (52) It may be that Fischer assumes that the 'can' in 1) and 2) entails 'not causally determined', but I see no reason why an account consistent with determinism cannot be given. So I think it's premature, at *this* stage in the discussion, to declare a "dialectical stalemate".

At the beginning of Chapter 5 and then again in Chapter 9, Fischer sketches something he calls "the Basic Version of the argument for Incompatibilism" and claims for it the following virtue: It articulates, *without relying on counterfactuals*, the intuitive ideas behind the Ginet/van Inwagen argument, avoiding the stalemate by providing a helpful restructuring of the problem. He first presents the main premise as: "Someone can do x only if his doing x can be an extension of the actual past, holding the laws fixed." (88) Later he tells us that we should read this as equivalent to: "Someone can do x at t only if there exists a possible world with the same past relative to t and the same natural laws as in the actual world and otherwise "suitably related" to the actual world in which S does x at t." (91)

As the main premise in an argument for incompatibilism, this is disappointing. It's a logical consequence of the truth of determinism that there are *no* worlds with the same past and laws where anyone ever does anything other than what she in fact does. To accept Fischer's premise is to accept the view that determinism is incompatible with the freedom to do otherwise. This premise "restructures" the debate only in the sense that it's a clear and concise statement of the incompatibilist thesis. We will make progress in the debate only insofar as Fischer gives us *reasons* for accepting this premise, reasons independent of the clashing intuitions about counterfactuals that he thinks were responsible for the "dialectical stalemate".

Fischer gives two reasons for accepting his incompatibilist premise. The first reason is that the premise articulates a "very natural, intuitive picture of the future as a garden of forking paths"; "at various nodes, there are several pathways that represent genuinely accessible futures, but there is *just one single path* that can be traced *backward* in time." (197) Fischer says that this picture represents how we

think about time and possibility; we think that we have genuine alternatives only insofar as we think that there are branching paths ahead of us. The second reason is that some of the assumptions of rational choice theory are best explained in terms of this branching future picture. Even if a so-called “backtracking” counterfactual is true (eg. even if it’s true that if Sam were to decide to go skating today, then there would have been no accident yesterday), “it is reasonable to restrict one’s attention to those possible worlds one can actualize...one can only actualize those possible worlds which are extensions of our world (in the sense that they share the past of the actual world)”. (96)

Let’s look at the rational choice argument first. I agree that when we deliberate it’s *natural* for us to restrict our attention to those actions we think are possible given both the past and the laws. But does rational choice-making *require* this restriction? Suppose that determinism is true. If so, then there is only *one* extension of the actual past, given the laws. Does Fischer think that then there is no such thing as rational choice? I assume not. So the fact that we reason in accordance with Fischer’s branching future assumption—as if determinism is false, as if we have libertarian free will—is irrelevant. It gives us no reason to believe that we *should* reason this way or that we have alternatives *only if* determinism is false and we have libertarian free will. At best, it gives us reason to believe that we should reason as if determinism were *nearly* false and as if we have the next best thing to libertarian free will—that is, by restricting our attention to worlds where our choice is the divergence miracle from past history.

This leaves the branching paths picture. I’m not sure I understand how this picture is supposed to provide an *argument* for the incompatibilist premise. A picture is not an argument. At best, it’s an intuition pump, something that could serve as the starting point of an argument. For instance, the picture of ourselves as puppets, moving in response to strings pulled by the puppeteer, counts as part of an argument for incompatibilism only insofar as the story is spelled out (eg. does the puppeteer leave our minds free, or does he pull neural strings as well?) and we are told in what relevant respects a deterministic agent is supposed to be like a puppet (eg. is determinism supposed to entail that none of our mental states are genuine causes?).

An intuition pump for incompatibilism is a story that gives us reason to believe that someone may act for reasons without having any genuine alternatives (perhaps because her reasons are not causally efficacious, perhaps because there are underlying physical factors which are symmetrically overdetermining causes of what she does). The problem with the branching path picture is that it doesn’t give us any reason to believe this. The contrast is supposed to be between climbing up a mountain on a path with many branches (indeterminism, libertarian free will, genuine alternatives) and climbing up a mountain on a nonbranching path (hard determinism, no genuine alternatives). So insofar as there is an argument based on an intuition pump (“if determinism is true, it’s just like...”), the relevant picture is the picture of climbing up a mountain on a path with no branches. The problem, though, is that while the linear path story is one in which we have fewer options, it’s not a story where alternatives are excluded altogether. Even if we

cannot get off the path we are on, we can stop and smell the roses or we can keep walking, we can walk slowly or fast, we can sing or talk philosophy. And so on. If this is what life under determinism is like, it’s not so bad; there is still lots of “elbow room” left. But of course the incompatibilist doesn’t think that this is how it would be under determinism; the incompatibilist thinks that life under determinism would be more like being strapped onto a seat on an out-of-control trolley or like being a puppet.

Fischer might reply that this is not the way he intended the picture to be understood. Each path represents a *single* action or sequence of actions; the branch points represent the *only* points at which there are choices between alternative actions. So we shouldn’t take the picture literally, as if there were further alternatives *within* a path.

But if that is the point of the story, then it’s no longer an intuition pump for incompatibilism; it is a pictorial representation of the incompatibilist thesis that someone can do otherwise only if there is a world with the same past and laws where she does otherwise. But what we were looking for was a *reason* for believing the incompatibilist thesis.

Near the end of the book, Fischer suggests that the real point of the branching path picture is not to serve as an argument but as “an enormously helpful heuristic device” (197) which helps us “see what is at stake” in the debate in a way that’s an improvement on the Conditional argument. “...it can help us to see the structure of the issues—the logical terrain—more clearly”. (199) But I don’t think this is an accurate depiction of how Fischer in fact uses the picture. He uses the picture, *selectively*, as an intuition pump in a way that’s intended to make us think that an apparently possible future action is a genuine alternative for us only if we are “able to get there from here”, that is, only if “one could in principle ‘draw a line’ that connects this possible future pathway with the pathway that represents the actual past.” (198) And in seeking to persuade us of this, he relies on our beliefs about actual mountain paths, including our beliefs about what *counterfactuals* are true: For instance, he tells a story in which we have just spent five hours walking up a mountain path. There is another path, not far away, but it’s separated from us by a deep gorge. “To get there we would have to have started out on that different path at seven in the morning.” (97) This intuition pump succeeds; it seems obvious that we cannot “get there from here”, that getting to the other path is no longer an option for us. But the intuition pump succeeds because we believe that a backtracking counterfactual is true; we think that we *cannot* get to the other path because we think that if we did, the *past would have to have been different*. It looks like Fischer is relying on our assumptions about the relation between a counterfactual and a ‘can’ claim in order to persuade us to accept a picture which commits us to incompatibilism. But if that is so, then Fischer’s “Basic argument” for incompatibilism isn’t in any way more basic, or less controversial, than the Conditional argument.

I don’t think we should be persuaded by Fischer’s argument. And I’ve argued that the Conditional argument doesn’t present a compelling case for incompatibilism either. Nevertheless, I am sympathetic to Fischer’s concerns. Libertarian-

ism (and thus incompatibilism) comes naturally to us; compatibilism is hard work. So let's ask what's at stake here. What difference would it make if hard determinism were true—that is, if we were never able to do anything other than what we in fact do?

Fischer's answer is surprising, even given the fact that he's been convinced by Frankfurt's argument that we may be responsible for our actions even if we can never do otherwise. He says that if he discovered tomorrow that determinism is true, then he would feel "irritation" about having gone through life with a very basic false belief (the belief that he has alternatives), but he would not feel disturbed in any deep way; he would not experience "the sort of profound disappointment and resentment" that he would if he discovered he was a brain in a vat. (214)

But it's hard to believe we would react so sanguinely, even if we think that the contrast between libertarianism and hard determinism is represented by the contrast between the branching and nonbranching path picture. Suppose we've spent our entire life thinking that there are many branching paths in our future, and now we come to believe that there is only one path, which never branches. Maybe Fischer is right in thinking that we would continue to care deeply about how we walk down our single path, take pride in what we do, and think of ourselves as responsible for our actions. But it seems that we would have to make some very fundamental changes in at least two respects. We would have to change how we think of ourselves as agents in relation to the future; our deeply engrained commonsense view is that we deliberate and choose among acts that are *genuine alternatives*. And we would have to change the way we think about moral responsibility for our omissions, for those acts we describe as "lettings", "permissings", and "allowings".

Fischer acknowledges that there is work to be done in defending an account of practical reasoning and responsibility for omissions that is consistent with the metaphysical beliefs of the hard determinist. (206) But I think he underestimates how deeply this will change our ordinary way of thinking about these things. Our commonsense view holds that there's a moral difference between doing and allowing harm, but also says that we are responsible for at least some of the harms we allow to happen. But this view rests on an assumption about alternatives: We don't hold a person responsible failing to prevent a harm unless we think that she *could have prevented it*; we don't even describe an act as an "allowing" of harm unless we believe that the person could have prevented the harm. If hard determinism is true, then this belief is *always* false. So far as practical reasoning is concerned, it seems that we believe, and perhaps can't help believing, that we make choices among acts each of which is such that we are *able to do it*, given the circumstances that in fact obtain. So learning the truth of hard determinism means learning that all of our practical reasoning has been, and perhaps will unavoidably continue to be, based on a false belief. But it's even worse than that for Fischer. Some incompatibilists think that the question of the truth or falsity of determinism is irrelevant to questions of rational choice; that is, they think that we can

figure out whether it's rational to do x rather than y without settling the question of whether determinism is true or false. But Fischer's discussion of Newcomb's problem suggests that he believes that whether it's rational to do x rather than y (take one box or two) may depend on whether or not we have alternatives, and thus depend on whether or not hard determinism is true.

Fischer defends two main claims in his book—he defends the view that determinism is incompatible with being able to do otherwise, and he defends Frankfurt's way of arguing that we may be responsible for what we do even if we can never do otherwise. Although I'm not convinced by his arguments for either claim, I found this book very much worth reading. Among its strengths are its clear presentation of some highly technical and complex arguments, particularly the various versions of the Ginet/van Inwagen argument for incompatibilism. Some very useful points are made along the way—eg. Fischer's arguments against the thesis he dubs "Restrictivism". But if I'm right in thinking that Fischer's heroic attempts to defend Frankfurt are, in the end, unsuccessful, then it matters, much more than Fischer thinks, whether determinism is compatible with the ability to do otherwise. And while I disagree with Fischer's claim that the debate between compatibilists and incompatibilists has reached a "dialectical stalemate", I share his respect for the Ginet/van Inwagen argument as a sophisticated starting point for thinking about the problem. Where does this leave us? I think it leaves us right in the middle of those hard metaphysical and modal questions that a generation of Frankfurt-influenced compatibilists have sought to avoid.

Notes

¹ For the modal argument, see Ginet 1990, and van Inwagen 1983. See also Lamb 1977, for an early version of the argument. For Frankfurt's argument, see Frankfurt 1969.

² Frankfurt's article continues to generate an enormous literature. For just a small sampling, see Blumenfeld 1971, Clarke 1994, Fischer and Ravizza 1991, Ginet 1996, Lamb 1993, McIntyre 1994, Widerker 1995.

³ The question of responsibility for omissions is more complex. Some philosophers, including Clarke 1994, Frankfurt 1988, and McIntyre 1994, think that there are Frankfurt-style counterexamples for omissions as well as actions. Others, including van Inwagen 1983 and Fischer and Ravizza 1991, disagree, holding that someone is responsible for failing to perform an action only if she could have performed that action.

⁴ Ginet 1996 is an exception.

⁵ All numerals in parentheses refer to page numbers in Fischer 1994.

⁶ See, for instance, Lehrer 1968.

⁷ A libertarian might not concede that this is possible, but a compatibilist should.

⁸ See Lamb 1993.

⁹ Widerker 1995 makes a similar claim, but he presents his argument as based on "the libertarian conception of freedom" (p.259). My argument is not so restricted; my claim is that no matter what your conception of freedom to do otherwise, Frankfurt stories fail.

¹⁰ See Nagel 1979.

¹¹ See Fischer 1983.

¹² See Lewis 1981.

¹³ For the articulation of one such argument, and a compatibilist reply, see Vihvelin 1991.

¹⁴ See van Inwagen 1989, Dennett 1984 at p.133, and Wolf 1980 at pp. 152–153.