

## BOOK REVIEWS

*Responsibility and Control: A Theory of Moral Responsibility*, by John Martin Fischer and Mark Ravizza, S.J. Cambridge, UK: Cambridge University Press, 1998. Pp. viii and 269 + 7 page index. \$59.95 (hardcover). \$19.95 (paper).

JOHN J. DAVENPORT, Fordham University

*Responsibility and Control* is today the leading work in the compatibilist tradition. Its argument that determinism is compatible with responsibility for actions, omissions, and their consequences is the comprehensive statement of a theory that began with Fischer's 1982 article on "Responsibility and Control."<sup>1</sup> Mark Ravizza joined this project in his first joint paper with Fischer on responsibility for inevitable actions and events.<sup>2</sup> The main purpose of this new volume is to solve problems (highlighted by Peter van Inwagen, Carl Ginet, David Zimmerman, Eleonore Stump, Randolph Clarke, and others) with Fischer's recent sketch in *The Metaphysics of Free Will* of the theory of moral responsibility in terms of "guidance control" (which does not require the power to bring about alternatives).<sup>3</sup> After a brief overview of *Responsibility and Control*, I will consider a new problem for the use of overdetermination cases in arguing for compatibilism. Finally I will show how this problem applies to the authors' new analysis of the origin of moral responsibility.

1. *An overview of the book*

In chapter one, the authors introduce their largely Strawsonian conception of moral responsibility and explain the direct and indirect threats determinism can pose to moral responsibility. They focus on the "freedom-relevant condition" of moral responsibility, i.e. the kind of *control* it requires rather than epistemic conditions associated with voluntary action (p.13). Following Harry Frankfurt's original critique of the Principle of Alternate Possibilities (PAP),<sup>4</sup> chapter two presents the basic argument from cases involving counterfactual and simultaneous interveners to "semicompatibilism" or the thesis that moral responsibility does not in general require the kind of libertarian freedom or "regulative control" that is indeed incompatible with global causal determinism. This basic argument is now

well-known to many readers. Consider a familiar case such as "Assassin," in which Sam shoots the mayor on his own, but could not have done otherwise because Jack had previously implanted a radio-controlled neurostimulator into Sam's brain and was ready to use it to force Sam to shoot the mayor if he showed any sign of wavering (*Responsibility and Control*, p.29). Since Sam is still responsible for shooting the mayor, the authors regard such cases as evidence that the agent's moral responsibility derives from features of the *actual sequence* that caused the action, rather than from alternative possibilities, even when these are available (p.37). The "guidance control" that responsible agents must display in the actual sequence has two aspects: it requires that the intentional process or psychological "mechanism" leading to the action be *same*, or responsive in certain ways to reasons for action, and that this same psychic process be recognized by the agent as "his own," or as an expression of his agency (p.38-39).

In chapters two and three, the authors concentrate on improving their account of the kind of reasons-responsiveness required by the first aspect of guidance control. They argue that moral responsibility requires that the mechanism from which the action flows (a) be "moderately" responsive to reasons (meaning that the agent can recognize a suitable range of reasons for doing otherwise, including moral reasons, that form a consistent and intelligible pattern) and (b) be at least weakly reactive to such reasons (meaning that the agent would do otherwise at least some of the time she recognizes such reasons). This account solves many of the problems with the earlier model of guidance control in terms of "weak reasons-responsiveness," such as the unwanted implication that agents who are only erratically responsive to bizarre reasons can be responsible. Problems remain, which I cannot explore in this short review, but this new account presents one of the most sophisticated compatibilist models of sanity available in contemporary moral psychology.

In chapter four, the authors are also largely successful in presenting a new account of "guidance control" over consequences of our acts and omissions, which provides a very effective response to van Inwagen's past arguments that responsibility for consequence-universals or states of affairs intuitively requires alternative possible consequences in this sense. Fischer and Ravizza show that in a range of cases, including van Inwagen's examples, what determines whether or not the agent is responsible for the consequence is not whether it is avoidable or inevitable, but whether the *agent's act* could have made a difference to the consequence in the prevailing environment (pp.101-112). To test this, we hold fixed the background circumstances and omit other "triggering events" that count as causes of the relevant consequence at the same time or after the agent's act causes it (pp.113-120). Here again some problems remain, particularly in distinguishing between triggers and background conditions. But the onus will be on incompatibilists to find any serious flaws with the plausible position the authors map out.

The analysis of responsibility for omissions in chapter five includes innovations based on this treatment of consequences.<sup>5</sup> The authors begin with a "wide conception of omissions" according to which all failures are omissions,<sup>6</sup> and then canvass two sets of cases: in the first, it seems like the



unavoidability of an omission defeats responsibility: in the second, Frankfurt<sup>7</sup> they argue that responsibility for simple failures to act is just a special case of responsibility for the bodily actions that constitute such failures (p.133). They then analyze complex omissions as cases in which the omitted bodily action brings about "a relatively narrowly specified negative consequence universal" (p.134). Just as for other consequences, they argue that when the unavoidability of this consequence-universal seems to defeat responsibility for the complex omission, this is not in fact because of its unavoidability *per se* but rather because its absence is not dependent in the required way on the absence of the agent's bodily action (p.135). This new analysis yields a symmetrical treatment of responsibility for actions and omissions, contra Susan Wolf's defense of an asymmetrical approach (pp.148-150).<sup>8</sup>

In chapter six, Fischer and Ravizza also respond impressively to van Inwagen's so-called "direct argument" for the incompatibility of determinism and responsibility by showing how simultaneous and counterfactual intervener examples can be devised to undermine van Inwagen's key transfer principle in the direct argument. This analysis is impressive, but I doubt it can yield a positive argument for responsibility-determinism compatibilism, because it cannot show that inevitability is compatible with responsibility when the "ensuring condition" operates in the agent's *own path* within the actual sequence as a whole, as of course it would if universal determinism held. But this Achilles heel of Fischer and Ravizza's response is a reflection, I think, of a deeper difficulty about the use of overdetermination cases to argue for compatibilism—the subject I turn to now.

## II. The fundamental problem with overdetermination-cases arguments against PAP

Several different types of objection have been raised against the kind of preemptive and simultaneous overdetermination cases on which Fischer and Ravizza rely throughout their book. These problems can be distinguished by looking at the general structure of overdetermination cases. Remember that the function of such examples is supposed to be to invalidate alternative-possibility conditions on responsibility by building our intuitions that a given person *P* can be morally responsible for something *X* (a decision, an intention, a bodily action, or the consequences of such an action) even when *X* is inevitable or (equivalently) when possible worlds in which *X* does not occur are inaccessible to *P*.

The persuasiveness of every overdetermination example, whether it involves simultaneous or counterfactual interveners, depends on a *division of labor* between two elements: one (*R*) which explains why the agent is responsible for *X*, and another (*I*) which explains why *X* is inevitable. The division consists in the fact that *R* and *I* are not in the same "pathway" (in Fischer and Ravizza's sense), even if both are present somewhere in the actual course of events as a whole. This anatomy of overdetermination cases can be illustrated first with the simple sort of example that started the disputes. In Frankfurt's "Neurologist" case, the inevitability-explaining

factor (*I*) is Black, who is prepared to intervene by triggering electrodes in Jones's brain to make him vote for Carter if he demonstrates any likelihood of trying to vote for Reagan. *X* of course is Jones's voting for Carter. And *R* is apparently (in the unanalyzed form of the example) just the fact that Jones deliberated about politics, formed the intention to vote for Carter, walked into the ballot box, and pulled the lever registering his vote all on his own, without any manipulation. Of course the largest part of the compatibilist's task is to analyze in much subtler terms just what conditions must be met by this process in the "pathway" leading up to Jones's actions for it to make *X* imputable to him. All the Neurologist example tries to do is show that an ability to bring it about that he refrains from *X* is no part of the *R* that makes *X* imputable to Jones. This point requires that the *R* and *I* factors do their work separately: the features of the scenario that make *R* imputable and the features that make *X* inevitable must be *different and independent* features.

We can now distinguish three broad types of objection to such cases. (1) Given this general structure, "Neurologist" and similar examples leave it up to further argument to show that the relevant *X* for which the agent is responsible and the *X* that is inevitable are the *same X*. This is the main issue in the dispute about whether Frankfurt-type cases leave "flickers of freedom."<sup>9</sup> In chapter four on "Responsibility for Consequences" the authors respond effectively to van Inwagen's argument that we are not responsible for event-universals we cannot avoid, but only for event-particulars that we can avoid despite the presence of preemptive or simultaneous interveners. This still leaves open questions about other flicker-defenses of PAP, but since these are not addressed in *Responsibility and Control*, I will leave them aside here.<sup>10</sup>

(2) The structure of overdetermination cases also leaves it up to further argument to show that the *I*-factor in such examples ensures the inevitability of *X* without tacitly presupposing a kind of determinism inconsistent with libertarian or agent-causal conceptions of freedom.<sup>11</sup> This is the issue in the ongoing debate about whether the signs that "trigger" intervention by counterfactual interveners imply a question-begging determinism. Fischer and Ravizza do not pursue this debate in their book, but Fischer and others have defended the authors' position in recent exchanges on this question.<sup>12</sup>

(3) There is a third and I think more fundamental kind of objection that focuses on the *R*-element in overdetermination cases that are supposed to count against PAP-principles. Fischer and Ravizza have done an admirable job addressing objections of type (1) and (2) by isolating a single set of *R*-features or responsibility-ensuring conditions that may explain responsibility for the relevant *X*s across a variety of challenging cases. But although Fischer and Ravizza use overdetermination examples with great acuity, such examples can *at best* show that responsibility for some *X* can sometimes be *locally* compatible with *X*'s unavoidability within a restricted context of inquiry that limits the focus to the relevant *X*. The very nature of overdetermination cases prevents them from showing us that the agent's responsibility for the relevant *X* is ever *ultimately* compatible with determinism.

Here's why. Such cases will intuitively persuade us that some *X* is both

imputable and unavoidable *only if we assume that the agent is responsible for R*, that is, for the feature (however interpreted or filled out) that makes his responsibility for X intuitive. But this assumption will seem plausible to all parties in the dispute only if they do not simultaneously stipulate that R itself is inevitable for the agent. Nor can the same overdetermination example focusing on X show us that R's inevitability would be compatible with the agent's responsibility for R, as required for his responsibility for X. Overdetermination cases instead rely on our presupposing that the agent's responsibility for R is plausible, without inquiring into the conditions of this responsibility, for their intuition-building work even to begin. For instance, in "Neurologist," we have to accept that Jones is responsible for deciding and intending to vote for Carter, before the example can demonstrate that responsibility and inevitability are locally compatible in the case of Jones's voting for Carter.

The same holds even if we start from a different compatibilist account of R in overdetermination cases. Take Frankfurt's willing addict, who in Frankfurt's view is clearly responsible for taking heroin because he "identifies" with his desire for the drug (and with its compulsiveness). Here R is explained in terms of the agent's higher-order (or authoritative, or deep-self-expressive) will to be an addict.<sup>13</sup> X = the agent's taking the drug, and I = the psychophysical addiction or the literal irresistibility of the desire. The factors that make X inevitable and imputable operate separately, and the example demonstrates that the agent is responsible for an inevitable X only if we presuppose that he is already responsible for the higher-order will that confers responsibility on X. This presupposition is not (and could not be) demonstrated nor its conditions analyzed in the willing addict case itself.<sup>14</sup>

We find the same in more complex cases concerning responsibility for inevitable consequences. Take "Joint Assassin," in which Jack and Sam simultaneously shoot the mayor. Here the X for which Sam is apparently responsible is "the consequence universal, *that the mayor is shot*" (p.117). Fischer and Ravizza explain Sam's responsibility for this in a different way than Frankfurt did with his addict: in their analysis, R is a complex structural feature of the situation, namely that Sam's act issued from a suitably reasons-responsive mechanism and controlled the relevant outcome in (roughly) the sense that in the given causal environment, minus other triggering events in other pathways (such as Jack's shooting), X would not have come about had Sam acted otherwise. On this account, R is a feature of the *agent's path* in the actual sequence: it specifically abstracts from the pathways in which other simultaneous triggers operate in the total actual sequence leading up to the mayor's untimely demise. But the I-factor that makes X inevitable is just such a trigger, namely Jack's firing his gun when and how he did. Now R and I do their work separately to ensure that X is both imputable to Sam and unavoidable. Yet R makes Sam responsible for X only if Sam is already responsible for R. And Joint Assassins cannot show us that his responsibility for the factors that constitute R (on the authors' analysis) is itself compatible with the inevitability of R.

Every successful overdetermination example can be decomposed in the same way, and the same will be found in every case: the persuasiveness of such examples presupposes that we accept the agent's responsibility for

whatever feature R stands behind our intuition that he is responsible for the inevitable X, and this in turn requires that the R and I factors work independently. Compatibilist analyses of such examples have generally focused on giving an account of precisely what features play the R-role or explain our intuition that the agent is responsible for the relevant X; they have not shown that responsibility for the factors they pick out as R is compatible with R's inevitability. Nor could they use overdetermination examples for this purpose without falling into a vicious regress.

For instance, suppose Fischer and Ravizza wanted to show us that Sam's guidance control over X (the mayor's death) in Joint Assassins would still make Sam responsible for X even if we stipulated that the features of the actual sequence constituting this guidance control were also inevitable for Sam. So they add some other factor F to Joint Assassins, such that F makes R (the guidance control) itself inevitable. Then F would by definition operate through the same pathway as R, not in a different pathway as the I-factor does in persuasive overdetermination cases. As a result, suddenly the background assumption that Sam is responsible for R—which seemed so innocuous as long as the focus was on his responsibility for the mayor's death (=X)—becomes doubtful, because we have departed from the regular structure of overdetermination cases which alone preserves the intuitive unison of responsibility and inevitability. We would need some further reason not supplied in this 'Modified Joint Assassins' case to show why it is still plausible that Sam is responsible for X now that we stipulate that the moderately reasons-responsive mechanism leading to his decision to shoot was itself inevitable for Sam.

Suppose the authors tried to supply such a reason by inventing a different case (C) in which the features that play the role of R in Joint Assassins now become the relevant X. Then in C we would have to have a new R and I for this X, and the agent's responsibility for our new R would again be presupposed and unexplained in C. There is no evident escape from this regress for the compatibilist. Thus in any overdetermination case that preserves the plausibility of the agent's responsibility for X, the possibility will be left open that what explains the agent's responsibility for R is that R (or some part of it) was agent-caused, or that R (or some part of it) was elected from among multiple alternatives the agent had the power to bring about. And either way, this means that the agent's responsibility for X will not be *ultimately* compatible with determinism. Nothing will prevent libertarians from explaining responsibility for an inevitable X in any given overdetermination case by saying that the agent satisfied libertarian conditions in bringing about (some part of) the features that play the R-role in that overdetermination case. This amounts to subsuming such cases via libertarian tracing principles.

Thus *even if* the worries about triggers and flickers of freedom can be overcome, Fischer and Ravizza's examples will only succeed in building part of the intuition the authors want: they may indeed show that agents can sometimes be responsible for unavoidable event-universals, acts, and perhaps some decisions, and hence that responsibility and inevitability are sometimes *locally* compatible. But libertarians are already prepared to concede this much on the strength of traditional examples: a man who

knowingly locks himself in a room is responsible for remaining therein, even though he cannot do otherwise, because his predicament is a foreseen result of an unmanipulated choice between alternatives. Libertarians modify their PAP-type principles with *tracing principles* to cover such cases, just as Fischer and Ravizza allow that an agent can be responsible in a derivative manner for an action that does not result from a moderately reasons-responsive mechanism if this itself is due to prior acts for which she is responsible in the primary manner, because they flowed from moderately reasons-responsive mechanisms (p.49-50). Responsibility for actions is thus locally compatible with actions flowing from irrational mechanisms on Fischer and Ravizza's account, but not *ultimately* compatible with non-reasons-responsiveness. Libertarians make analogous accommodations in tracing principle revisions to PAP principles for decisions, intentions, actions, and consequences. Thus if overdetermination examples show no more than a local compatibility of responsibility and inevitability, then they are no serious threat to libertarian intuitions about the ultimate conditions on moral responsibility. Therefore such a "tracing-defense" against overdetermination cases is not *ad hoc*: it focuses on a problem that any theory of responsibility must address, whether it is compatibilist or incompatibilist.<sup>15</sup>

### III. The New Account of Responsibility for Psychological Mechanisms

Since this tracing objection was presented in such a schematic form, it will help to see how it applies more concretely to Fischer and Ravizza's own positive account. In *Responsibility and Control*, the authors confront the problem that a suitably reasons-responsive mechanism M can make us responsible for an action, omission, or their consequences, only if we are already responsible for M. Perhaps, the incompatibilist will say, we are only responsible for M if it (or some crucial part of it) is agent-caused or chosen from among alternatives we had the power to realize.<sup>16</sup> Indeed such agent-causal and libertarian accounts of responsibility for psychological mechanisms provide one clear way out of a familiar objection to Fischer's earlier statement of his model, i.e. that suitably reasons-responsive mechanisms could be produced in the agent by brainwashing, neurological manipulation, divine intervention, and so on. Libertarian or agent-causal accounts would deal decisively with this *source-problem* by ensuring that we are not responsible for psychological mechanisms deriving from coercion or artificial interference. Fischer and Ravizza therefore had to describe and defend an alternative compatibilist account of responsibility for psychological mechanisms.

They do this in the final two chapters of *Responsibility and Control*, which lay out an impressive argument for an "historical" conception of responsibility according to which the agent acquires the subjective conditions for regarding himself as a moral agent (or a fair target of reactive attitudes) through a process of acculturation and initiation into the language-game of holding persons responsible and ascribing responsibility to oneself. Through three stages — "training," "taking responsibility," and "being held responsible" by others (p.210) — the young person learns to become responsible for the sort of psychological mechanisms that can in turn make her

responsible for particular acts, omissions, and consequences. The crucial step is the middle one: "The process by which an agent takes responsibility for the springs of his action makes them *his own* in an important sense" (p.210). Although I cannot do justice to the details of their account, in sum taking responsibility for mechanisms involves three basic conditions:

- (a) First, the individual learns to see her psychological states — such as belief, desire, choice, and intention — as "the causal source...of upshots in the world" (p.210-11).
- (b) Second, "the individual must accept that he is the fair target of reactive attitudes as a result of how he exercises this agency in certain contexts." If the individual is not prone to philosophical reflection, this need involve no more than a sense that it is "part of our given social practices" for others to praise, blame, and react to him as a result of his agency in some non-arbitrary range of contexts (p.211).
- (c) Third, the individual's internal view of himself as a moral agent constituted by these first two sets of beliefs must be based in "an appropriate way" on "his evidence for these beliefs," which will include recognizing the causal effectiveness of his choices and actions, learning the moral language game by experiencing social practices (p.213), and presumably learning that he can conform his actions to social expectations. The training and moral address of the community constituted in this language-game provides the necessary evidence for the individual's self-ascription of responsibility to count as rational (p.214).

The result of this process is that the individual becomes accustomed to thinking of actions that flow from familiar kinds of psychological mechanisms (including not only reflective deliberation and prudence but also habit and unreflective desire-gratification) as flowing from *her*. Acceptance of action-ownership through mechanism-ownership becomes virtually automatic unless there are special features of the case that could lead the agent to exempt herself. Thus we do not take responsibility directly or case-by-case for each particular mechanism-taken behind each of our acts: rather, "Having taken responsibility for behavior that issues from a kind of mechanism, it is almost as if the agent had some sort of standing policy with respect to that kind of mechanism" (p.216).

On the basis of this theory, Fischer and Ravizza argue that an agent cannot take responsibility for certain psychological mechanisms that are not reasons-responsive, such as irresistible desires or subconscious suggestion implanted by neurological manipulation (p.232). Similarly, although a reasons-responsive mechanism could be directly induced by external interference, agents have normally taken responsibility for such mechanisms only as "uninfluenced by such factors as subliminal advertising, direct electronic stimulation of the brain, and so forth" (p.233). The complete manipulation-mechanism will not be among those mechanisms for which

the agent has taken responsibility, since even if she thinks she is acting on a normal psychological mechanism, she is not: "In a case in which a moderately reasons-responsive mechanism is implanted without an agent's awareness, the agent has not taken responsibility for the kind of mechanism that issues in action" (p.234).

Finally, the authors consider whether an agent's standing policy of recognizing that she is ordinarily responsible for actions flowing from familiar kinds of psychological mechanisms could itself be manipulated. They admit, "it is conceivable that the individual's view of himself as an agent and an apt candidate for the reactive attitudes be electronically implanted" (pp.235-6). But they respond that such an agent would not have formed the required view of himself "in the appropriate way," based on reality-tracking evidence, as required by the third condition on taking responsibility (p.236). Instead he would have been deceived, deluded, or tricked, and should not be held responsible by one who knows the full story.

Despite its improvements, problems remain in this account. Drawing on Robert Heinlen, Eleonore Stump has argued that an agent could be manipulated into taking responsibility for a psychological mechanism M that he knows has been implanted (perhaps by inculcating a false philosophical ideology).<sup>1</sup> Her case seems to be analogous to Gettier-cases in epistemology: the beliefs that constitute "taking responsibility" for M on Fischer and Ravizza's account are not directly implanted, but are rather motivated by staging the kind of appropriate evidence the agent needs to be warranted in such beliefs, on their account. Such cases show that an omnipotent manipulator could bring it about that the agent has true and justified beliefs about his psychological mechanism that constitute "taking responsibility" for it, yet not really be responsible for it. This also suggests how hard it is to trace responsibility for actions to responsibility for mechanisms without tracing mechanism-responsibility itself to incompatibilist sources that rule out manipulation.

But I want to focus on a more limited point: it is at any rate impossible to use overdetermination cases to argue that the conditions of responsibility for psychological mechanisms are ultimately compatible with determinism. In chapter eight §7, the authors argue that counterfactual-intervener cases give a philosophically reflective agent reason to believe that responsibility for psychological mechanisms is compatible with their being inevitable. For in such cases:

the actual-sequence mechanism is, intuitively, the agent's own, whereas the alternative-scenario mechanism is not. But the Frankfurt-type cases are entirely consistent with causal determinism's being true. Another way of making the point is that causal determinism in itself does not imply that all sequences are relevantly similar to the alternative scenario of a Frankfurt-type case, in which there is, for example, significant direct manipulation of the brain (and thus the mechanism is not the agent's own) (p.228).

This is true, but it misleadingly implies that absence of direct interference from other agents, machines, or natural events must be sufficient for the

mechanism to count as the agent's own. Counterfactual-intervener cases do not show this: they merely presuppose that the mechanism on which the agent acts in the actual sequence (as opposed to the counterfactual sequence) is her own, and invite the reader to go along with this assumption for the purpose of seeing that the agent can thereby become responsible for some inevitable decision, action, or consequence. We cannot learn from such cases whether the agent's ownership of (or responsibility for) the mechanism on which he acted requires that the mechanism originated in libertarian choice, or agent-causation, or neither. Thus if a Frankfurt-type case is one in which the agent is by hypothesis responsible for what she does 'on her own' in the actual sequence, it has yet to be proven that "Frankfurt-type cases are entirely consistent with causal determinism's being true," as the authors assume.

This is the deepest problem with Fischer and Ravizza's argument, and it is relevant to a wide body of literature in which it has not been clearly recognized before. But this should not obscure the fact that *Responsibility and Control* is full of challenging new ideas and ingenious responses to better-known objections. Aside from being well-written and enjoyable to read, it is the fruition of the entire tradition stemming from Frankfurt's 1969 paper, which has become the dominant tradition in compatibilism today. Although some amendments may be proposed, I doubt that a more systematic presentation of this tradition's approach will soon be attempted or achieved. This makes the book essential reading not only for graduate seminars on free will but also for philosophers interested in related problems about freedom and divine foreknowledge.

#### NOTES

1. Fischer, "Responsibility and Control," *Journal of Philosophy*, 89 (January 1982): 24-40, reprinted in John M. Fischer, ed., *Moral Responsibility* (Ithaca, NY: Cornell University Press, 1986): 174-190. Also see John Fischer, "Responsiveness and Moral Responsibility," in *Responsibility, Character, and the Emotions*, ed. Ferdinand Schoeman (Cambridge: Cambridge University Press, 1987): 81-106.
2. Fischer and Ravizza, "Responsibility and Inevitability," *Ethics* 101 (1991): 258-278. In addition, see the authors' joint introduction to Fischer and Ravizza, eds., *Perspectives on Moral Responsibility* (Ithaca, NY: Cornell University Press, 1993); Fischer and Ravizza, "Responsibility for Consequences," new in *Perspectives on Moral Responsibility*: 322-347 [also appearing in Jules Coleman and Allan Buchanan, eds., *In Harm's Way: Essays in Honor of Joel Feinberg* (New York, NY: Cambridge University Press, 1994)]; and Fischer and Ravizza, "Responsibility and History," in French, Uehling, and Wetstein, eds., *Midwest Studies in Philosophy 4: Philosophical Naturalism* (Notre Dame, IN: University of Notre Dame Press) 430-451. These (and other) articles introduce many of the main ideas developed and synthesized in *Responsibility and Control*.
3. John Fischer, *The Metaphysics of Free Will* (Cambridge MA: Blackwell, 1994).
4. Harry Frankfurt, "Alternative Possibilities and Moral Responsibility," *Journal of Philosophy*, 66:23 (December, 1969); reprinted in *The Importance of what*

the care about (Cambridge University Press, 1988): 1-10, and also reprinted in *Moral Responsibility*, ed. Fischer: 143-152.

5. Some of these were presented in John Fischer's paper, "Responsibility, Control, and Omissions," *Journal of Ethics* 1 (1997): 45-64.

6. They argue against narrower alternatives by suggesting that in cases where other passive or active factors in the scenario would prevent John from saving a drowning child if he had tried, "John omits to save the child, and yet in neither case does he have the power to save the child." So in overdetermined cases, "the agents seem to omit to do certain things they in fact cannot do" (p.125, notes 3 and 4). But I find this unpersuasive: John does have the power *not* to omit saving the child. Had he tried fruitlessly to save him, competent speakers would say "John failed to save the child" but not "John omitted to save the child." Trying and failing to do X is neither doing X nor omitting to do X. Since he could have tried, he could have avoided the omission, even though he could not have done X. Libertarians are bound to emphasize this feature of the authors' examples.

7. See Harry Frankfurt, "What we are morally responsible for," in *How Many Questions? Essays in Honor of Sidney Morgenbesser*, ed. L.S. Cauman, Isaac Levi, Charles Parsons, and Robert Schwartz (Indianapolis: Hackett Publishing, 1983); reprinted in *The Importance of What We Care About*, 95-103, esp. pp.101-102.

8. The authors consider a different account favored by the asymmetrist, according to which complex omissions are constituted by "an agent's *not causing* the relevant positive state of affairs," instead of bringing about a negative state of affairs (p.146). On this analysis, it follows that the agent cannot have guidance control of the positive state of affairs not caused, but the authors argue that what matters instead is guidance control of "the omissions themselves" rather than the positive state of affairs they avoid causing (p.149). This is intriguing, but it is not entirely clear that *not causing* is something over which the agent could have guidance control without having regulative control either.

9. This phrase was introduced by John Fischer in his book *The Metaphysics of Free Will* (pp.134-147).

10. There are various flicker defenses of PAP, and some focus on other residual possibilities aside from the possibility of some different event-particular occurring than the one actually brought about by the agent's action. These strategies are not discussed in *Responsibility and Control*, but Fischer treats them more thoroughly in a forthcoming paper in *The Journal of Ethics* and in "Recent Work on Moral Responsibility," *Ethics* 110.1 (October, 1999): 93-139. A full analysis of these other flicker-objections, which I hope to attempt later, will show that some of them are simply picking up on symptoms of the deeper problem I identify in this section.

11. Contra recent trends, I reserve "libertarianism" for conceptions holding that moral responsibility requires the power to bring about or actualize alternative possibilities (of various kinds pertaining to different contexts of responsibility) and thus I distinguish between "libertarian" and "agent-causal" theories of free action. Although some of its proponents assume that agent-causation only occurs in the process of forming one volition out of a range of possible freedoms open to the agent, which makes agent-causation a form of libertarian freedom in my sense, the concept of agent-causation is compatible with lacking volitional access to such robust alternatives. I do not speak of theories requiring only agent-causation, without the ability to agent-cause otherwise, as "libertarianism without alternative possibilities" (since I think this apparent contradiction in terms only generates unnecessary confusion). Thus in my terms, an agent-causal theory of the freedom required for moral responsibility may be

incompatible with determinism without thereby counting as "libertarian."

12. See David Widerker, "Libertarian Freedom and the Avoidability of Decisions," *Faith and Philosophy* 12.1 (January 1995): 113-118; John M. Fischer, "Libertarianism and Avoidability: A Reply to Widerker," *Ibid.*: 119-125; David Widerker, "Libertarianism and Frankfurt's Attack on the Principle of Alternative Possibilities," *The Philosophical Review* 104.2 (April 1995): 247-261; David P. Hunt, "Frankfurt-Counterexamples: Some Comments on the Widerker-Fischer Debate," *Faith and Philosophy* 13.3 (July 1996): 395-401; David P. Hunt and Charlotte Katzoff, "Avoidability and Libertarianism," *Ibid.*: 415-421; Alfred Mele and David Robb, "Rescuing Frankfurt-Style Cases," *The Philosophical Review* 107.1 (January 1998): 97-111; and now Daniel J. Speak, "Fischer and Avoidability: A Reply to Widerker and Katzoff," *Faith and Philosophy* 16.2 (April 1999): 239-247.

13. At least identification with the desire that motivated the act is sufficient for responsibility for the act in Frankfurt's view. His treatment of other examples implies that it may not be necessary, however. As I read him, the necessary condition Frankfurt isolates in other cases is one of omission: that the agent did *not* identify with any desires opposed to the desire on which she acted. See Frankfurt, "Freedom of the Will and the Concept of a Person," *Journal of Philosophy*, 68.1 (January, 1971); reprinted in *The Importance of What We Care About*: 11-25.

14. This leaves it open for libertarians to argue that among the conditions on responsibility for higher-order willing is a libertarian requirement of being able to avoid a given volitional identification or being able to identify otherwise. In that case, the willing addict will seem responsible for taking the drug, even though he could not avoid taking it, only because we tacitly assume that it was open to him in the past to be an unwilling addict. I explored this idea in "Conditions for Freedom of the Higher-Order Will," Frankfurt and Augustine," presented at the Eastern Division meeting of the American Philosophical Association, (Atlanta, GA, December 28, 1996).

15. This argument may not seem to cover the sort of "global" Frankfurt-type cases developed by Mele and Robb in "Rescuing Frankfurt-style Cases," *The Philosophical Review* 197.1 (January 1998), by David Hunt in "Moral Responsibility and Unavoidable Action," *Philosophical Studies* 96 (December 1999), and by Eleonore Stump in "Alternative Possibilities and Responsibility: The Flicker of Freedom," *Journal of Ethics* (1999): 1-26. I believe my account will handle such cases, but this will require further argument.

16. And then clearly, even if the action, omission, or consequence were inevitable, our responsibility for it would again still ultimately have incompatibilist conditions.

17. Stump, "Persons: Identification and Freedom," *Philosophical Topics* 24.2 (Fall 1996): 200-206.

*The Conscious Mind: In Search of a Fundamental Theory*, by David Chalmers

JOHN DUPRÉ, Birkbeck College, University of London and University of Exeter

The Conscious Mind is certainly an ambitious book, and in many ways an impressive one. Chalmers argues for a bold and unfashionable thesis. The argument is often ingenious and ambitious, and at the (many) points