BOOK REVIEW

John Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge University Press, Cambridge, 1998, 277 pp., GBP37.50 (hardback).

John Fischer and Mark Ravizza's book *Responsibility and Control: A Theory of Moral Responsibility* is an attempt to develop a unified account of moral responsibility for actions, omissions and consequences that takes seriously "Frankfurt examples" and their many variations. (Roughly, a "Frankfurt-example" is one in which an agent acts voluntarily, and seems responsible for her action, despite the fact that some unactualized force waits in the wings to ensure that she so act should she begin to act differently.[1]) The book is not an unqualified success, but it is a *qualified* success for it both presents a theory of moral responsibility that anyone working in the field should examine – even if only to reject – and brings out the various relevant factors involved in numerous Frankfurt-style examples with impressive clarity and perspicuity.

The view of moral responsibility presented develops a distinction between two different types of control – regulative control and guidance control – discussed in Fischer's earlier book *The Metaphysics of Free Will*.[2] An agent has regulative control over her action when she has alternative actions available to her. Agents can have guidance control even when they have no alternative possibilities available to them. Fischer and Ravizza claim that guidance control is what is required for moral responsibility,[3] and argue, also, that guidance control is compatible with causal determinism. This means, as Fischer and Ravizza stress, that moral responsibility is derived from what *actually happens* in the causal sequence leading to an agent's action, not what could have happened, but didn't. Much of the book is spent in presentation, and argument in favor, of their view of guidance control.

Their view of guidance control develops the idea, suggested by Fischer and Ravizza in earlier work, that the crucial capacities exercised by morally responsible agents are capacities for recognizing and responding to reasons for action. To be a morally responsible agent, according to

Fischer and Ravizza, is to guide one's conduct in accordance with those features of oneself or one's environment which provide rational justification, at least to oneself, for the conduct in which one engages. It is a powerful and appealing idea, well worthy of development. However, I have some reservations about the way in which the idea is developed. My first worry concerns the promise of Fischer and Ravizza's theory to account for the responsibility-undermining force of physical constraints; the second is methodological: the authors have not succeeded in really grounding their view in a conception of the point or purpose of ethical evaluation.

Physical Constraints

Any theory of moral responsibility should tell us why it is that physical constraints undermine the kind of freedom required for moral responsibility. Oddly enough, Fischer and Ravizza don't directly address this issue, although they do discuss a variety of cases – such as one in which sharks would prevent a man from saving a child were he to try to – in which physical constraints appear to be playing some kind of role. However, it is quite difficult for Fischer and Ravizza to provide a satisfactory explanation for the occasional responsibility-undermining force of such constraints. As Fischer and Ravizza are committed to the claim that responsibility derives from features of the actual causal sequence through which an agent's action comes about, they are committed to the claim that physical constraints are only responsibility-undermining when they play some role in the actual causal sequence, when, that is, they are part of the "mechanism" through which an agent's action comes about.

Certainly, there are cases in which, for instance, we struggle with bars and fail to break them, and, arguably, those constraints do play a role in the mechanism through which our actions come about in such cases; I stay in a room as a result, in part, of the fact that the door is locked when I struggle with it to no avail. But these are not the only cases in which the presence of physical constraints of various sorts accounts for our lack of responsibility for what we do or fail to do. For instance, imagine that I am in a room and quite rationally, and truly, believe a lock on the door to be strong enough to prevent me from leaving no matter what I do. As a result, I consider it fruitless to struggle with the door and decide not to waste the time and effort; I thus stay in the room without bothering even to try the door. Fischer and Ravizza's theory had better not say that I am morally responsible for staying in the room in such a case. The lock on the door prevents my leaving, and the lock plays a role in the actual mechanism through which my staying comes about: it plays a role in bringing about my belief that I am locked in.[4] And, although I am not sure if Fischer and

Ravizza would develop the point this way, their theory does say that I am not responsible in this case, as it should. I explain.

For Fischer and Ravizza, guidance control is exhibited when the mechanism through which an agent's action comes about is "moderately reasons-responsive" and is "her own". Both of these conditions are given further analysis. One of the features of moderate reasons-responsiveness is "weak reactivity" to reasons: there must be some possible world in which the mechanism through which the agent actually acted is held fixed, but in which the agent has sufficient reason to act differently from the way she did act, does act differently, and acts differently for that reason. Weak reactivity to reasons is, then, a necessary, but insufficient condition on guidance control. And, it seems in the example just described, the mechanism through which my staying in the room comes about is not weakly reactive to reasons: there is no possible world (appropriately similar to the actual world) in which the mechanism is held fixed – and thus the lock is held fixed – but in which I have sufficient reason to leave *and leave*. So long as the lock on the door is strong enough to prevent my leaving, I'm staying, no matter what reasons I have to leave. So far, so good.

But consider a second case in which all is just as in the first case with one exception: unbeknownst to me, the lock is weak from corrosion and would be broken if I were only to give the door a good push. I still believe, quite rationally and for no reasons for which I am to be faulted, that the lock is too strong to be broken, so I still make no effort to force the door, decide to stay in the room and stay as a result of my decision. But, there is a possible world in which I have sufficient reason to leave and leave: say I observe some atrocity outside that I have so much reason to prevent that the mere observation of it induces panic in me: despite my cool and rational belief that there is nothing that I can do to get out of my confines, I throw myself at the door, and, 'lo and behold, I exit the room, for the lock, remember, is corroded. In this non-actual possible world, I have sufficient reason to leave the room, and I leave the room for that reason. What's the result? Fischer and Ravizza are committed to saying that I am responsible for staying in the second case (assuming, as we may, that I satisfy the additional conditions on guidance control) and not in the first. But the only difference in the two cases is corrosion of a lock my ignorance of which is no fault of my own, nor indicative of culpable failure of rationality on my part. How can that make the difference in moral responsibility?

Notice that this objection does not turn on any claim to the effect that I am responsible in the first case and not responsible in the second, or vice versa. It depends only on the claim that the two cases are *relevantly similar* with respect to responsibility. But then the two cases together present Fisc-

her and Ravizza with a dilemma: If I am responsible in both cases, then the first shows that weak reasons-reactivity – and thus guidance control as understood by Fischer and Ravizza – is not necessary for responsibility; and if I am not responsible in both cases, then the second serves to illustrate the insufficiency of guidance control for responsibility. Since it appears that weak reactivity is the best tool available to Fischer and Ravizza for accounting for the responsibility-undermining force of physical constraints, their theory simply doesn't provide an adequate account of the full range of factors that can serve to undermine our responsibility.

Fischer and Ravizza do not attend closely to the issue of mundane physical constraints, for they are concerned to explain the irrelevance to responsibility of features of ourselves or our environments that play no role in the actual mechanisms through which we act, and the most obvious such factors – obvious only because of the recent history of the philosophical literature on alternate possibilities – involve pernicious neuroscientists or un-triggered dispositions towards irresistible urges. But the fact remains: our responsibility can be undermined by ropes and chains, and a comprehensive theory of moral responsibility should explain why.

The Method of Reflective Equilibrium

Early in the book, Fischer and Ravizza say they aim to develop "reflective equilibrium" between intuitions with regard to particular cases and a theory of moral responsibility. In the first instance, any effort to gain reflective equilibrium must involve development of a theory that accounts for intuitions about cases. But it is just as important to be ready to sacrifice our intuitions about some particular case, or class of cases, out of a desire to adhere to the theory. What would justify this desire in the case of a theory of moral responsibility? One possibility is that we would want to hold on to a theory, even in the face of conflict with pre-theoretic intuition, because of some important conceptual connection between the theoretical account of moral responsibility and some conception of the point or purpose of morality, or our practices of assigning praise and blame.

For instance, P F. Strawson's essay, "Freedom and Resentment"[5] proposed a strategy for justifying accounts of moral responsibility: say why agents who have the features specified by the account are appropriate objects of the moral sentiments or, as he called them, the "reactive attitudes". Strawson, however, said more than just this, for he offered a gloss on the purpose of the moral sentiments relevant to his conception of moral responsibility: the moral sentiments, he claimed, are responsive to "quality of will" exhibited by others and ourselves in so far as, and to the degree that, we see those others as capable of meaningful participation in the kind

of special relationships with which human life is peppered. Hence, we are responsible for all and only those actions that are expressive of quality of will, when we are such agents. Strawson's account of moral responsibility, then, is justified through appeal to the function and role of the reactive attitudes.

Advocates of a theory of moral responsibility owe us a justification of the theory that plays the same role as that offered by Strawson. Such a justification must do more than simply confront the account with our intuitions with regard to particular examples. The justification is required even if the advocates of the theory intend to defend and develop the theory through a process aimed at reaching "reflective equilibrium". Without such a justification, we are at a loss to decide if the theory requires revision in the face of examples or not. Fischer and Ravizza, unfortunately, do not, in the end, offer any such justification; they simply move back and forth from theory to examples demonstrating, often with great ingenuity, ways in which the theory, despite first appearances, accounts for the examples, or ways in which the theory needs to be elaborated so as to account for some particular example.

It is true that Fischer and Ravizza see themselves in the tradition of which Strawson is a part. They describe their project as one of determining what it is that an agent must be like if she is to be an appropriate candidate for the reactive attitudes. However, while they do tell us what an agent needs to be like, in their view, for this to be the case – she must act from her own moderately reasons-responsive mechanism – they don't offer a conception of the point or purpose of the reactive attitudes which explains why those attitudes attach to all and only those agents who act from such mechanisms. Further, it is not clear that the conception that Strawson had in mind will serve the turn, for it is not clear whether or not the capacities to act from mechanisms of the sort described by Fischer and Ravizza is even necessary for engagement in the kind of meaningful relationships the support of which Strawson saw as the fundamental purpose of the reactive attitudes. This is a serious lacuna in their account which I hope they will fill in further work.

In short, then, Fischer and Ravizza should be applauded for their subtle and nuanced development and examination of "Frankfurt examples", but their book fails both to take seriously the relevance to moral responsibility of the mundane cases of physical constraint, and fails to develop connections between their view of moral responsibility and the point and purpose of moral evaluation.

NOTES

[1] Such examples were originally developed in Harry Frankfurt, "Alternate Possibilities and Moral Responsibility" in *The Importance of What We Care About*, pp. 1–10, Cambridge University Press, 1988.

[2] John Fischer, *The Metaphysics of Free Will*, Blackwell, 1994.

[3] They also make mention of "epistemic" conditions that an agent need satisfy to be responsible, but their scope does not include detailed discussion of these conditions.

[4] Fischer and Ravizza might deny this. They might claim that the lock itself is not part of the mechanism through which my action comes about. However, if they were to make this claim, it would not be clear how they could account for the responsibility-undermining force of the lock in the example under consideration without abandoning their adherence to an actual sequence view. That is, the natural thing to say is that the lock undermines responsibility because there are not relevant alternative sequences in which the agent acts in the manner that the lock prevents her from acting. Fischer and Ravizza, however, cannot say this without abandoning their belief that only features of the actual mechanism through which an agent's action comes about are relevant for assessing the agent's moral responsibility.

[5] Reprinted in Fischer and Ravizza (eds.), *Perspectives on Moral Responsibility*, Cornell University Press, 1993, pp. 45–66.

University of Southern California   GIDEON YAFFE