

constitutive. She takes such goods to be something like the typical effects of the practice, regardless of how they would be understood or valued by participants in that practice. As a result, her account seems to largely repackage familiar consequentialist objections to torture in more edifying Aristotelian language. *Mainstreaming Torture* does do valuable work in describing the realities of how torture has been practiced by the United States in recent years and the way it has deformed our political culture. However, Gordon's treatment of the philosophical issues falls prey to an endemic weakness of virtue ethics, the tendency to preach to the choir. *Mainstreaming Torture* helps remind us of much that we already know, but it offers little guidance with respect to issues about which we may be still making up our minds.

DAVID SUSSMAN

University of Illinois at Urbana-Champaign

Levy, Neil. *Consciousness and Moral Responsibility*.

Oxford: Oxford University Press, 2014. Pp. xii+157. \$45.00 (cloth).

Debates concerning the conditions under which agents are morally responsible for their actions typically focus on free will or control. Less commonly, discussions attend to the so-called epistemic condition or to issues concerning the quality of an agent's will. It's quite rare, however, that anyone will invoke consciousness (in any guise) as relevant to moral responsibility. After all, you might think, we're often self-deceived about what we're up to and so not conscious of the motives that are actually moving us to action. And yet, this doesn't obviously vitiate our status as morally responsible for those actions. For example, suppose that Jerry sincerely believes that he supports some discriminatory policy on the basis of a political principle. Given that such a policy is discriminatory (and, we might suppose, not grounded in any justified political principle), it seems that we can reasonably judge that Jerry is blameworthy for supporting it. And yet finding out that Jerry's sincere belief is false and that he really supports the policy because of unconscious racist attitudes (perhaps attitudes he knows nothing about) does not obviously undermine this judgment. Consciousness of one's motives, at the very least, cannot plausibly be a necessary condition on moral responsibility. And because similar considerations can be adduced against the claims that an agent must be conscious of the causes of her actions, their consequences, and so on, it seems that consciousness per se is not necessary for moral responsibility.

Yet despite the apparent plausibility of the above considerations, in *Consciousness and Moral Responsibility*, Neil Levy powerfully argues that a certain kind of consciousness is in fact a necessary condition on moral responsibility. Of course, even this statement of Levy's view isn't entirely perspicuous, since we'll need to know (i) what Levy means by "consciousness" and (ii) which facts need to be represented by consciousness in order for any agent to be responsible. Concerning the latter issue, Levy does not argue that to be morally responsible for an action we must be conscious of why we are performing that action, where this

entails an understanding of the whole array of mental and physical processes necessary for action. Instead, he argues for a weaker (and to my mind, a much more plausible) claim: to be morally responsible for an action, an agent must be conscious of the facts that give an action its moral significance. In specifying what this claim entails, and in particular, what he means by “conscious of,” Levy offers us an account of consciousness itself (or at least, the kind of consciousness that might be necessary for moral responsibility). Importantly, he does not insist that in order to be conscious of such facts, one must represent these facts in one’s “mind’s eye” at some time just before deciding what action to perform. Rather, on Levy’s view, one must be “aware” of such facts in the sense that these facts are integrated into one’s “global workspace,” which is the mechanism by which domain-specific considerations are processed and assimilated by the agent holistically. Indeed, this mechanism is precisely what Levy means by “awareness” or “consciousness,” and its proper functioning is what is putatively necessary for morally responsible action, since it is in virtue of an agent’s capacity to process and assimilate domain-specific considerations that her behavior and reasoning are flexible and suitably sensitive to morally significant features of her environment.

Levy claims that an agent is morally responsible only if she is aware of the facts that give her actions their moral significance. This thesis, the *Consciousness Thesis*, apparently stands in sharp opposition to recent empirical work by cognitive scientists and conceptual work due to philosophers of action. Concerning the empirical work, Levy does a good job of undermining the putatively threatening findings of Benjamin Libet and Daniel Wegner, arguing (convincingly) that not only do they fail to impugn the Consciousness Thesis, but they fail to show much of anything that might be relevant to moral responsibility. If there is a problem with the Consciousness Thesis, then, it does not seem to be with its empirical commitments.

Having dispatched empirical threats to the Consciousness Thesis, Levy turns to his philosophical opponents. Nomy Arpaly, for example, has argued that although Huck Finn is not aware of why it’s wrong to return Jim to the authorities, he is nevertheless praiseworthy (and morally responsible) for deciding to not rat him out (*Unprincipled Virtue: An Inquiry into Moral Agency* [Oxford: Oxford University Press, 2003]). Angela Smith has similarly argued that we need not be able to “consciously entertain” attitudes in order to be morally responsible for those attitudes (“Responsibility for Attitudes: Activity and Passivity in Mental Life,” *Ethics* 115 [2005]: 236–71). These views, it seems, are very much in line with the sorts of considerations that make it plausible that Jerry is blameworthy for supporting the discriminatory policy even though he is not aware of the fact that his support is really motivated by racist attitudes and not commitment to political principle. What unifies these opponents, Levy tells us, is that they “deny that information that plays a role in shaping behavior needs to be available for easy and direct access in order for agents to be morally responsible for that behavior” (35). And this, for Levy at least, seems to be a clear indication that they must reject the Consciousness Thesis.

But are these opponents really opponents of the Consciousness Thesis as Levy characterizes it? It’s not clear to me that they are. After all, according to Levy’s articulation of the exact content of the Consciousness Thesis, it seems consistent

with some fact being integrated into an agent's global workspace that it is not necessarily "available for *easy* and *direct* access" (emphasis added) for the agent. Moreover, when Arpaly claims that Huck Finn can be praiseworthy for refusing to return Jim to the authorities even though he is unconscious of the reason why that behavior is in fact the right thing to do, the sense of consciousness at stake is distinct from that Levy is claiming to be necessary for responsibility. So there seems to be no conflict between Arpaly's claims about Huck Finn and the Consciousness Thesis (at least the version of the Consciousness Thesis that Levy endorses). This is clear, for example, when she says, "Huckleberry is not capable of bringing to consciousness his unconscious awareness [that he shouldn't turn Jim in to the authorities]" (Arpaly, *Unprincipled Virtue*, 73). But for Arpaly, it's (in part) in virtue of Huck's "unconscious awareness" that he is praiseworthy, so we shouldn't read her as denying that some form of awareness is necessary for moral responsibility. Thus, what Arpaly says seems consistent with Levy's claim that consciousness is necessary for moral responsibility (and so, for praiseworthiness), since the notion of "unconscious awareness" Arpaly invokes is apparently no different from the notion of consciousness or awareness that Levy takes to be necessary for moral responsibility. No doubt, Arpaly describes the kind of awareness in question as "unconscious awareness," but it's clear that the state in question nevertheless involves integration of information into Huck's global workspace (to put the point in Levy's terms), since for Arpaly, it seems that the awareness in question is "unconscious" only in the sense that at the time of action, Huck is unable to articulate to himself the relevant moral facts (even though he is able to recognize and respond to them, and is hence "aware" of them). This suggests that in some ways, the dispute between Levy and his philosophical opponents is not as deep as he makes it out to be.

Of course, even if this is right, and Levy's opponents aren't actually opposed to the Consciousness Thesis when it's suitably spelled out, it doesn't show that the Consciousness Thesis isn't a philosophically important one. Nor does it show that it isn't true. So to further motivate the Consciousness Thesis, Levy considers the case of Kenneth Parks, who killed his mother-in-law. Parks's crime, however, wasn't premeditated. Nor did it appear to be a crime of passion. Instead, Parks was sleepwalking at the time at which he killed his mother-in-law, and as a result, he was subsequently found not guilty of murder in a Canadian court. We can make sense of this verdict, Levy thinks, by appealing to Consciousness Thesis. To develop this point, Levy draws on some of the recent work in neuroscience, which suggests that somnambulism is a disorder of consciousness. And since what went wrong in the case of Parks could be traced to a disorder of consciousness, it's plausible to conclude that a properly functioning ability to integrate domain-specific inputs (this is the role for consciousness, after all) is necessary for responsibility. In other words, the best explanation for why Parks should be excused for the killing is that his somnambulism rendered him unaware of the facts that made his actions so heinous, even a simple fact like that he was hurting someone. I found Levy compelling on this point, but it's worth pointing out that there is a more general explanation for why Parks should be excused that Levy doesn't consider.

This more general explanation of why Parks should be excused is simply that at the time of action he (or the mechanism that moved him to action) was

not suitably responsive to reasons to refrain from killing his mother-in-law. This sort of explanation has been favored by a number of theorists, including John Fischer and Mark Ravizza, Michael McKenna, Dana Nelkin, R. Jay Wallace, and Susan Wolf (to name a few). It's true that on its face, this account of why Parks should be excused seems to stand in sharp opposition to the explanation offered by Levy, since there is no invocation of consciousness by reasons-responsiveness theorists. But this is ultimately mistaken. For what the Consciousness Thesis requires of agents is that they be aware of the facts that give actions their moral significance in a way that allows for flexible response to those facts. But this is the precise reason why so many have taken "reasons-responsiveness" to be necessary for moral responsibility. It seems to me, then, that the Consciousness Thesis, rather than being a wholly independent condition on an agent's being morally responsible (as Levy sometimes presents it) is really a way of spelling out the idea that in order to be morally responsible for an action one must be able to respond to reasons (which are just facts, after all) for action in a flexible way. Seen in this light, Levy's development of the Consciousness Thesis is a success. For although he says less concerning the modal properties of a "responsive" agent than earlier accounts (e.g., Fischer and Ravizza's account), he actually provides a (partial) explanation of why something like reasons-responsiveness is required for moral responsibility: reasons-responsiveness, when understood as an awareness of the facts that give one's actions their moral significance, is necessary for sensitive and flexible action itself. And this point, which seems to follow from Levy's account of consciousness as a mechanism by which agents are able to integrate domain-specific information in a holistic way, is an important one.

Levy concludes his brief but deftly argued volume with two chapters devoted to arguing that the main theories of moral responsibility are really committed to the Consciousness Thesis. Given the connection between Levy's thesis and the claim that reasons-responsiveness is necessary for moral responsibility, it's not surprising that control theories of responsibility are committed to this thesis (this is discussed in chap. 6). Perhaps more surprising is Levy's claim (in chap. 5) that real-self theories of responsibility are committed to the Consciousness Thesis. But even this isn't too surprising given the account of consciousness Levy develops earlier in the book.

Neil Levy's *Consciousness and Moral Responsibility* is short, and it's a quick read. (In all the familiar ways, this is both a virtue and at times, a [minor] vice.) The view developed is very plausible, if not wholly novel, and in his defense of this view, Levy displays a great deal of philosophical sophistication. His discussion of the relevant empirical science might be old hat for some, but I suspect that many researchers working on these issues will find it extremely helpful as an introduction to the relevant literature (I certainly did). Moreover, he does push on some theories of moral responsibility (e.g., the real-self views) in interesting ways—ways that future researchers should pursue (e.g., what is required for an action to "express" an agent's real self?). Consequently, I suspect that the volume will be of significant interest to those working on issues in the philosophy of action and the philosophy of mind.

D. JUSTIN COATES
University of Houston