BOOK REVIEWS

Alfred R. Mele, *Effective Intentions: The Power of Conscious Will*. New York: Oxford University Press, 2009. ix +178 pp.

Al Mele's *Effective Intentions* is a much-needed antidote to some of the anti-realist views of free will that have been put forward recently by various scientists, most notably, Benjamin Libet and Daniel Wegner. Now, unlike Wegner, Libet actually wants to admit that humans have a limited sort of free will—a bit roughly, he thinks we have the freedom to consciously veto actions that are about to occur—but I think this is a pretty uninteresting kind of free will, and I'll gloss over it by focusing on the following hypothesis:

> (FW) Human beings have (non-veto-style) free will, that is, a kind of free will that goes beyond the limited veto-style free will that Libet has in mind.

Libet and Wegner both reject (FW), and they both do this for empirical reasons because they think (FW) is inconsistent with evidence that we currently have. Mele's book defends (FW) against their arguments. Moreover, it also defends the following against Libet- and Wegner-style arguments:

> (H) Our (overt intentional) actions are at least partially caused by our intentions (or the acquisition or persistence of our intentions, or the physical correlates of these things).

I think the arguments Mele provides in defense of (FW) and (H) are largely right. And I think his book is interesting, important, and well worth the read.

1. Summary of the Book

After laying out some conceptual background in chapter 1, Mele argues in chapter 2 that human beings can have unconscious intentions that are causally relevant to action. Now, to most contemporary philosophers, the idea that we have unconscious intentions probably sounds like a no-brainer, but psychologists and neuroscientists often write as if it's analytic that all intentions are conscious, so the point is worth arguing.

In chapters 3 and 4, Mele responds to Libet-style arguments against (FW) and (H). I discuss this at length below, so I won't go into it here.

Philosophical Review, Vol. 120, No. 3, 2011

In chapter 5, Mele defends (H) against Wegner-style arguments. The issues here center around Ouija-board-type phenomena, that is, phenomena in which people move their bodies without being aware that they're doing so (they think their bodies are being moved by spirits). Wegner thinks these cases generate problems for (H), but Mele disagrees. He begins by asking whether the people in question are *intentionally* moving their bodies. If they're not, then there's no threat here to (H) because there's no intentional action. And if they *are* intentionally moving their bodies, then it may be that they've got *intentions* to do this and that these intentions are causally relevant to their actions. Since intentions can be unconscious, this is clearly possible, and there's no evidence that it's not actual—that there aren't any causally efficacious intentions in these cases. (You might also try to argue against (H) by assuming that Ouija-board-type bodily movements aren't caused by intentions and inferring from this that bodily movements are in *general* not caused by intentions. But as Mele points out, there's no evidence for thinking that ordinary actions are caused in the same way that Ouija-board-type bodily movements are caused.)

In chapter 6, Mele takes up the topic of the accuracy of subjects' reports of the time of onset of conscious awareness of intentions to act. These reports play a crucial role in Libet's argument, but Mele argues that there are reasons to doubt the accuracy of these reports. Mele's arguments here are based on two facts, namely, (a) that the reports in question vary wildly from trial to trial, even for individual subjects, and (b) that the reported time of onset of conscious awareness of intention can be manipulated by stimulation to the brain after action.

After arguing in chapters 3–6 that extant scientific studies fail to undermine (H), Mele argues in chapter 7 that we have good empirical evidence for the positive claim that conscious intentions *are* causally relevant to our actions. He argues this point by appealing to studies about the effectiveness of *implementation intentions*. These studies suggest that if people form conscious intentions about *when* and *how* they're going to perform some action, then they're far more likely to perform the given action. Mele argues convincingly that this gives us good evidence for thinking that conscious intentions are at least sometimes among the causes of our actions. (It's worth noting that this thesis is much weaker than (H), which is a claim about *all* of our [overt intentional] actions. Mele defends (H) against objections, but he never provides a positive argument for it and never commits to its truth.)

Finally, in chapter 8, Mele describes some hypothetical discoveries that he thinks *would* undermine (FW) and (H).

I agree with most of Mele's arguments, but I'm not entirely satisfied with his responses to Libet. I turn to this now.

2. Libet, Mele on Libet, and Me on Mele on Libet

2.1. *Libet*

We've known since the 1960s that voluntary actions are preceded in the brain by a slow change in electrical potential known as the *readiness potential*, or RP. Libet set out to establish a time line for the RP, the onset of conscious intention to act, and the act itself. In his study—or, rather, one of his studies—subjects were instructed to flex their wrists whenever they felt like it; moreover, during the experiment, subjects were watching a clock, and they were asked to report the exact time that they became aware of an intention to flex. What Libet found, to put things very briefly, is that on average the RP appears about 550 ms before action (that is, before wrist flexing) and that subjects don't become aware of an intention to act until about 150 ms before action. From this, Libet infers that our actions are initiated unconsciously and, hence, that (FW) is false.

2.2. *Mele on Libet*

Mele wants to defend (FW) and (H) against Libet-style arguments. He makes three central points here. The first two have to do with a model that he proposes of what might be going on in the Libet study. The proposed model is as follows:

> *Mele's model*: (i) What arises 550 ms before action (that is, what's indicated by the RP) is a *potential cause of an intention to act* (for example, an urge, or a cause of an urge); (ii) the conscious intention arises about 150 ms before action (that is, at the time when, according to Libet, awareness of intention enters the picture); and (iii) this intention plays a causal role in producing the action, that is, the wrist flexing.

The first point Mele makes about this model is that it's perfectly consistent with all the data that we currently have, most notably Libet's data. Mele's idea here is that since Libet's data don't give us any reason to doubt the above model, they don't give us any reason to doubt (FW) or (H).

Mele's second point is that his model actually fits *better* with the data than Libet's model does. Libet's model can be defined as follows:

> *Libet's model*: (i) The RP is an indicator of an intention to act, which is generated nonconsciously by the brain about 550 ms before action, and (ii) the agent becomes aware of this intention around 150 ms before action.

Mele gives two arguments for thinking that his model is empirically more plausible than Libet's. They proceed roughly as follows:

> *Argument 1*: In one of Libet's studies, subjects were told to prepare to flex their wrists and then to veto the plan at the last second. These subjects

didn't exhibit ordinary RPs, but they did exhibit something very similar, which we can call *veto-RPs*. Now, we know that veto-RPs aren't indicators of intentions to flex because the subjects in these trials never had any intentions to flex (they were told at the start *not* to flex). Therefore, absent a reason to think that veto-RPs are indicators of completely different kinds of things from ordinary RPs, it's plausible to suppose that ordinary RPs aren't indicators of intentions either.

Argument 2: Libet's model suggests that it takes humans about 550 ms to move from intentions to actions (in cases like those in Libet's studies). But studies on human reaction times suggest that we're actually much faster than that.

Finally, the third point Mele makes in response to Libet's argument is that it doesn't even matter if Libet's model is true because that model is perfectly consistent with (FW) and (H). Indeed, Mele argues that (FW) and (H) are consistent not just with Libet's model but also with the following:

Nonmental Causation Model: (i) RPs are indicators of wholly nonmental brain processes that give rise to actions, and (ii) agents aren't aware that they're going to act until about 150 ms before action, and it's not until then that they have intentions to act.

According to Mele, the Nonmental Causation model poses no threat to (FW) or (H). As long as things like beliefs, desires, decisions, intentions, and so on are parts of the causal chains leading to actions, it doesn't matter if the causal processes begin before these mental things are activated. In short—and this is really the central claim Mele makes in response to Libet—it's not a problem for (FW) or (H) if it turns out that the causal processes that lead to our actions begin with brutely physical events that are wholly nonmental and nonconscious; this isn't a problem because the causal order could go like this:

wholly nonmental, nonconscious events → beliefs, desires, decisions, intentions, and so on → action

2.3. Me on Mele on Libet

Is Mele's response to Libet a *good* response? Well, that depends. It seems to me that what we should say in response to Libet depends on whether we're thinking of free will in a compatibilist or an incompatibilist way. If we're trying to defend compatibilist free will, then we can respond as Mele does. Indeed, it seems to me that Mele's response to Libet is exactly analogous to the classic compatibilist's response to the threat of determinism. In response to the worry that all of our actions might be causally determined by the big bang, Humean compatibilists respond with indifference, pointing out, first, that the causal order might go like this:

> the big bang → beliefs, desires, decisions, intentions, and so on → action;

and, second, that if this is indeed how things go, then we do have free will.

This compatibilist response to determinism is pretty transparently parallel to what I called Mele's third response to Libet. Mele doesn't draw out this parallel himself, but I think part of the strength of his response to Libet is the way it parallels the classical Humean response to the possibility of determinism. Indeed, if you think of free will in a compatibilist way, it's hard to see why Libet's studies would be worrying at all. If you're not bothered by the possibility that your choices and actions might be determined by the big bang, why on Earth would you be worried that they might be caused by nonmental, nonconscious brain events in your own head?

So far, so good. But I think more needs to be said in response to Libet, for I think his arguments are most naturally read as arguments against the idea that we have *libertarian* freedom. Now, to be sure, Mele's discussion succeeds in blocking *some* Libet-style antilibertarian arguments, but I think that the *best* Libet-style antilibertarian argument isn't touched at all by Mele's responses. To appreciate this, we first need to locate the relevant versions of libertarianism. It seems to me the most important decisions vis-à-vis the question of whether we have libertarian free will are what might be called *torn decisions*—that is, decisions in which the person in question has reasons for multiple options, feels torn as to which option is best, and decides without resolving the conflict, that is, decides while feeling torn. Given this, I think that the best versions of libertarianism are committed to the following thesis:

> (*) At least some of our torn decisions are undetermined in such a way that (a) the objective, moment-of-choice probabilities of the various tied-for-best options being chosen are all roughly even, and (b) which tied-for-best option is chosen is not significantly causally influenced (at the moment of choice) by anything external to the conscious choice itself, so that the conscious choice *is* the event that settles which tied-for-best option is chosen.

I don't have the space to argue here that (*)-style libertarian views are the best versions of libertarianism, but regardless of whether this is true, it seems to me that Libet's anti-free-will argument is best thought of as an argument against views of this kind. One might put the argument here as follows:

(1) Libet's studies suggest that conscious decisions (in particular, torn decisions) are causally influenced by nonconscious brain

processes that begin before we have the conscious experience of choosing. Therefore,

(2) Our torn decisions are not undetermined in the manner of (*).

Mele's discussion does nothing to block arguments like this. He's mainly concerned with defending the thesis that our actions are caused by our beliefs, desires, decisions, intentions, and so on. But Libet's argument isn't really an argument against that thesis; it's more naturally read along the lines of the argument in (1)–(2)—that is, as an argument for the antilibertarian claim that our conscious decisions are causally influenced by nonconscious brain processes in freedom-undermining ways. Thus, while I think Mele's responses to Libet are important for the reasons laid out above, and while I think compatibilists can rest content with Mele's responses, I also think that if we want a really complete response to Libet, we need to say more.

I should also say that I think there *is* a good libertarian response to the argument in (1)–(2). In a nutshell, what needs to be argued here is that there is no evidence for thinking that in torn decisions, the nonconscious brain processes involved in RPs are causally relevant to which tied-for-best option is chosen. I think this point can be argued, but I can't do it here.

Mark Balaguer
California State University, Los Angeles

Philosophical Review, Vol. 120, No. 3, 2011
DOI 10.1215/0031808-1263701

Michael Slote, *Moral Sentimentalism*.
Oxford: Oxford University Press, 2010. ix + 184 pp.

Self-described in the first sentence of its preface as "a more thoroughgoing, a more systematic, defense of moral sentimentalism than anything that has appeared since Hume's *Treatise of Human Nature*," Michael Slote's *Moral Sentimentalism* seeks to show how to use empathy as a common thread to tie answers to metaethical and normative questions into a single, broadly "sentimentalist" picture. We are never told precisely what it means for a view to count as "sentimentalist," but it's clear that Hume's own view is supposed to be a paradigm, and despite a number of differences, it's clear that Slote's own view is much indebted to Hume. As Slote points out several times, through an idiosyncrasy of etymology, Hume actually meant *empathy* by 'sympathy', and so in tying his view closely to the workings of empathy, Slote is hitching his wagon to the same horse as did Hume.

Beginning with an introductory discussion of the nature and effects of empathy, Slote articulates accounts of moral approval and disapproval by reference to empathy. The central part of the book then uses these accounts of moral approval and disapproval to address metaethical questions. In what is presented as one of the main accomplishments of the book, Slote first presents a "non-Kripkean" account of the "reference-fixing" of moral terms by appeal to his accounts of moral approval and disapproval. Then he moves on to explain why his account allows for a bridging of the "is-ought" gap and to explain why his account yields a more plausible account of moral motivation than its chief competitors. Finally, in the last part of the book, Slote tries to show how his empathy-based moral sentimentalism can yield plausible answers to normative questions about paternalism and particularly about justice. Its scope thus matches its aspirations to thoroughgoingness and systematicity.

So what is empathy? Rather than telling us exactly, Slote tells us a number of things about it. Our intuitive grip on empathy, to begin with, is supposed to come from the distinction between what Bill Clinton meant when he said, "I feel your pain," on the one hand (which is empathy) and feeling *for* someone who is in pain, on the other (which is sympathy). But whereas Hume seemed to think of this process as a simple kind of contagion, Slote takes seriously the idea that empathy is a real psychological capacity that individuals can possess to different degrees, and which can be more or less developed. Indeed, in contrast to Hume, the difference in empathy exhibited by different individuals turns out to be important for a number of Slote's other views.

But despite this important difference from Hume, Slote also emphasizes similarities. In particular, Slote sees Hume as having taken empathy ("sympathy") to play an analogous role in the moral realm as causation plays in the natural—a sort of "cement of the moral universe." Slote takes this to mean, at least in part, that empathy, like causality for Hume, works through spatial and temporal contiguity, which is why we have greater empathic reactions toward those near (and similar) to us than toward those far (and dissimilar) from us. Here, Slote believes, Hume was on to something important. The fact that empathy operates by contiguity plays a key role in allowing Slote to explain many of the broad contours of commonsense morality—the self-other asymmetry, the putative permissibility for middle-class Westerners of failing to make personal sacrifices in order to save impoverished lives in the Third World, the existence of agent-centered constraints that don't amount to absolute side constraints because they can be overridden if enough is at stake.

One question Slote never seems to have answered, however, is whether someone whose empathic identification spreads further than another always counts as *more empathetic*, or whether instead there is some point at which identifying better with those who are dissimilar from oneself or further away in space or time actually starts counting as *less* empathetic. It's fine to observe that empathy is a process that always involves identifying more closely with the near and similar and less closely with the far and dissimilar, but if other things equal it always counts as *more empathetic* or *at least as empathetic* to expand the circle outward (as seems initially highly plausible), then failures to seriously regard the interests of those very far away in space and time may count as a failure in empathy—at least a shortcoming relative to a more completely empathetic state. Since it appears to be central to Slote's account to allow that some kinds of action that involve paying less regard to those who are distant or dissimilar do not involve any shortfall in empathy, this would appear to commit him to the view that at some point identifying more closely with those who are dissimilar or distant from oneself does *not* make one count as more empathetic—a surprising conclusion.

To get moral judgment out of empathy, Slote takes an indirect strategy, through *moral approval* and *moral disapproval*. In contrast to Hume's *patient-based* account of moral approval, according to which it is what we feel toward an action or trait of character insofar as empathy leads us to approve of it from an impartial standpoint by leading us to identify with the feelings of the people whom it *affects*, Slote advocates an *agent-based* account of moral approval, according to which it is the approval that we feel for the agent of some action or trait of character insofar as we empathically identify *with the agent's state of empathy*. So rather than empathizing with the people who might be helped or harmed by some action, a fully developed empathic capacity is supposed to lead us to identify with *empathetic people* and hence approve of them (feeling "warmth" toward them), and to disapprove (feel "left cold" by) people who lack empathy.

It's not at all clear, however, how these feelings of warmth and coldness could simply be the result of a fully developed empathy. Of course, if our empathy is fully developed, then that will be one important similarity between us and other empathetic people, and hence should make us empathize more with them, other things being equal, because of that similarity. But empathy itself shouldn't lead us to feel "warmed" by their actions, unless they are themselves warmed by their own actions—after all, it is our capacity to feel their feelings, like they do, as Bill Clinton "feels your pain."

Even more problematically, it is very hard to see how it could be empathy itself that makes us feel "cooled" by the actions of those who are not empathetic. For if they are *unlike* ourselves, then other things being equal, we should empathize *less* with them. So all that would seem to follow from the nature of empathy is that we would react toward them in the same way that we react toward those who are distant from us in time or space—that is, to be unmoved. But Slote's account requires instead that something *positive* happens when we fail to empathize with those who are unempathetic—they must make us feel "cooled." So clearly there is something more going on here than empathy, which seems to straightforwardly leave the possibility that your empathy could be fully developed without your being able to experience moral approval and disapproval in this way.

In this brief review I've been raising concerns about the foundation for Slote's views—his accounts of empathy and of moral approval and disapproval. There is far more in the book, however—*Moral Sentimentalism* is wide ranging in its pursuits and resourceful in its quest to put empathy at the center of the answer to every question. It's no *Treatise of Human Nature*, but those who can avoid getting sentimental over what it's not will find much to learn from it.

Mark Schroeder
University of Southern California

Philosophical Review, Vol. 120, No. 3, 2011
DOI 10.1215/00318108-1263710

Keith DeRose, *The Case for Contextualism: Knowledge, Skepticism, and Context*. vol. 1. New York: Oxford University Press, 2009. xiii +288 pp.

It is hard to convey the philosophical excitement that was generated by attributer contextualism (henceforth "contextualism") in the mid-to-late 1990s. Although the basic nuts and bolts of the proposal weren't new—Stewart Cohen had, of course, been championing such a position for some time by this point—things really took off in a major way for this research program in the mid-1990s. There were two papers that changed the landscape in this dramatic way. One of them was David Lewis's "Elusive Knowledge," which came out in 1996. The other was Keith DeRose's "Solving the Skeptical Problem," which came out in 1995. Both are now classics of contemporary epistemology.

A key part in the tremendous wave of interest that ensued off the back of these two articles was the prospect that contextualism could *both* pay due attention to our actual practices of ascribing and denying knowledge while at the same time deal with fundamental epistemological problems, such as, in particular, the problem of radical skepticism (a central concern of both articles). Our hunch is that contextualism would not have generated nearly as much interest if either of these features of the view had been absent. This is important since both aspects of the contextualist research program have been roundly criticized in recent years.

We are grateful to Keith DeRose for helpful feedback on an earlier version of this review.

On the one hand, it has been argued that contextualism does not offer us the resolutions we seek for major epistemological problems. In particular, the contextualist response to the problem of radical skepticism has been attacked, with some claiming that, for example, it merely sidesteps the problem rather than resolves it. On the other hand, there has also been a concerted attack on the examples of (putatively) ordinary conversation that contextualism cites in order to motivate the position, such as the now ubiquitous "bank" cases. For example, experimental epistemologists have disputed the "data" in question, while proponents of the rival subject sensitive invariantist view (SSI) have offered very different explanations of what is going in the relevant cases (and also, of course, often disputed the "data" as well).

Arguably, then, contextualism faces a crisis point. If it is to dust itself off in the face of this critical onslaught and once more present itself as a credible philosophical proposal, then it needs to be able to offer a systematic defense of its main claims, one that is sensitive to the recent lines of critique. Hence DeRose's two-volume defense of contextualism, of which this volume is the first (the second is forthcoming), could not have come at a better time. The focus of this volume is on the ordinary language basis for contextualism, while the focus of the second volume is on the extent to which contextualism can help us with the big epistemological questions, especially with regard to the problem of radical skepticism. The main thesis of volume 1 is that contextualism provides the best account of our everyday practices of knowledge ascription and denial.

The Case for Contextualism (henceforth "TCC") consists largely of material drawn from a number of previously published papers. However, some new material has been included, most notably the entire final chapter, and an attempt has been made to indicate where the argument of older papers stands in relation to the current literature. For example, in chapter 3, previously published as "Assertion, Knowledge and Context," it is made clear that the argument of that paper doesn't aim to establish (as was claimed in the original version) that invariantists can't accommodate intuitions about bank cases used in arguments for contextualism as well as the knowledge norm of assertion (KAA). Instead, all that is claimed is that *classical invariantists* (that is, non-subject-sensitive invariantists) can't accommodate these intuitions while endorsing KAA. Taken as a whole, TCC presents an argument for its main thesis that deals with a large number of recent much-discussed objections to contextualism, especially those presented in recent work by John Hawthorne and Jason Stanley.[1]

Here's an overview of the argument of TCC. DeRose begins by providing what he takes to be the main motivation for adopting contextualism, that the way in which ordinary speakers use 'knows' in a number of everyday contexts

1. See especially, John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004) and Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2005).

(call these contextualist cases) prima facie supports contextualism (chapter 1), and offers a number of rules for how to best construct and present such cases in order to elicit procontextualist intuitions (chapter 2). A typical classical invariantist response to contextualist cases is to accept that varying standards govern whether speakers can *appropriately* ascribe knowledge to subjects of certain propositions but to argue that this doesn't reflect any variation in the standards for *truthfully* ascribing knowledge (this is known as a "warranted assertibility maneuver," or "WAM"). In chapter 3 DeRose argues that to make this response to contextualist cases the classical invariantist has to utilize a certain sort of WAM. In high-standards contexts one can't properly assert that one knows that *p*, whereas in low-standards contexts one can, just because in high-standards contexts one can't properly assert that *p*, whereas in low-standards contexts one can properly assert that *p* (that is, the classical invariantist must claim that standards for proper assertion vary with context). But if classical invariantists also adopt KAA, then they must think that the truth-conditions of 'I know that *p*' vary with context just because the assertibility conditions of '*p*' vary with context. So classical invariantists have two options. They can either deny KAA or not utilize this WAM.

Chapter 4 is an explanation of what DeRose calls the *gap view*. In any conversation the participants will have their own epistemic standards, so there is a temptation to think that the contextualist holds that any ascriber's knowledge ascriptions get their content from that ascriber's standards (what DeRose calls their *personally indicated content*). On the gap view, this is false. In general, in a conversation between two conversational participants A and B where the personally indicated content of A and B differ, an ascription or denial of knowledge to *S* is true/false iff *S* meets/fails to meet the standards set by the personally indicated content of both A and B, and *without a truth-value* iff *S* meets/fails to meet one set of standards but not the other.

Chapter 5 contains DeRose's response to a number of objections to contextualism based on linguistic data. Chapter 6 argues that, first, SSI is committed to the truth of strange sentences such as '*S* knows that *p* today, but on Tuesday, when the stakes will be higher, *S* won't know that *p*'. Second, it argues that, while contextualism is committed to the view that similar strange sentences ('*S* doesn't know that *p*, but what I said earlier, "*S* knows that *p*," was true') can be truthfully asserted, the same goes for views on which 'tall' is context sensitive ('*S* isn't tall, but what I said earlier, "*S* is tall," was true'), so there's no problem unique to contextualism.

The key battleground between SSI and contextualism is third-person knowledge ascriptions. In chapter 7 DeRose argues, first, that the SSI projectivist account of what goes on in High Ascriber-Normal Subject cases is insufficiently motivated and, second, that the contextualist can easily handle cases where it looks as if it should be the *subject's* epistemic standards that "call the shots." DeRose also argues that the contextualist has no problem accommodating the sort of data cited in defense of KAA and the analog knowledge norm for practical reasoning.

We'll briefly comment upon some limitations of the argument of TCC before considering objections. While DeRose acknowledges the question of why our intuitions about knowledge ascriptions and denials in contextualist cases carry such evidential weight (49–51), he doesn't offer much by way of response and, further, the growing body of literature in experimental philosophy suggesting that what he takes to be our intuitions about contextualist cases aren't shared by "the folk" isn't addressed.[2] The argument of chapter 3 is easily dealt with by denying KAA, and there's a growing body of literature suggesting that it should be rejected.[3] There is little discussion of skepticism and contextualist resolutions of skeptical arguments but, as DeRose makes clear, TCC is intended as the first of two volumes, the second of which will deal primarily with skepticism (44).

We will now discuss some objections to the arguments summarized above. First, DeRose's discussion of the gap view is unclear about a crucial detail. What he says is that, in conversational contexts in which A and B differ in personally indicated content, if A ascribes/denies knowledge that *p* to *S*, then that ascription/denial lacks a truth-value iff *S* meets/fails to meets the standards of one of but not both A and B. Does that mean that A's ascription/denial has an indeterminate semantic content? It would seem that this has to be so, but DeRose's discussion doesn't make that clear.

Second, DeRose wants to apply his gap view to the following sort of case. Say that in a low-standards context Ted asserts that Dougal knows that the bank is open on Saturdays. Later, Jack, who is in a high-standards context and was told about Ted's earlier assertion, says 'Rubbish, Dougal knows no such thing!'. DeRose says that Jack's assertion has no truth-value iff Dougal meets Ted's standards but doesn't meet Jack's standards, and so Jack's assertion has no truth-value (even though Ted's assertion was true, and what he asserted is still true in Jack's high-standards context). Consider this consequence. Later, Ted, still in a low-standards context, asserts that Dougal knows that the bank is open on Saturdays. Presumably on DeRose's view this assertion is without a truth-value iff Dougal meets Ted's standards but not Jack's, and this is the case. The upshot of the gap view seems to be that, for a subject who has been truthfully ascribed knowledge of a given proposition, if that subject is denied knowledge in a later

2. See, for example, Joshua May, Walter Sinnott-Armstrong, Jay Hull, and Aaron Zimmerman, "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study," *Review of Philosophy and Psychology* 1 (2010): 265–73. To be fair to DeRose, he does now have a paper forthcoming that addresses these issues, "Contextualism, Contrastivism, and X-Phi Surveys," *Philosophical Studies* (forthcoming).

3. For a useful survey of some of the main critical lines with regard to KAA, see Matthew Weiner, "Norms of Assertion," *Philosophy Compass* 2 (2007): 187–95.

high-standards context and then reascribed knowledge in a low-standards context, both that denial and reascription are without a truth-value. This is certainly not an innocuous consequence of the gap view.

Third, consider DeRose's objection to the SSI account of what goes on in High Ascriber-Normal Subject cases (chapter 7). Stanley has argued that what goes on in such cases is that ascribers incorrectly *project* their epistemic standards on to the subjects of their knowledge ascriptions. When we are inquiring into whether a subject knows that *p*, we are treating that subject as a *source of information*. This leads us to ask what that subject would know if he or she had our interests. There's an obvious problem, highlighted by DeRose. What about cases where an ascriber knows that *p* and is aware that this is so? The ascriber isn't going to treat the subject as a source of information because he or she already has that information. DeRose's response is that maybe the defender of SSI could provide some additional mechanism that explains why ascribers would still project their standards in such cases. But what the defender of SSI really needs is some sort of reason for thinking that this projection of epistemic standards is *mistaken*. It isn't enough to give a characterization of our epistemic practice that explains why we project our epistemic standards. The contextualist also thinks that we project our epistemic standards. The difference is that the contextualist thinks this is entirely correct. The characterization has to explain why we are mistaken in doing so.

If DeRose leaves his argument against projectivist SSI strategies here, he leaves himself open to the following sort of argument. The defender of SSI just has to give some reason for thinking that projecting our epistemic standards is mistaken. The response to this might just be: if there are good objections to contextualism, and the defender of SSI obviously thinks that there are, those objections should be reason enough.

One might think that the above is a real problem for DeRose. But it would be a mistake to think that DeRose is in such a dialectically weak situation. Unfortunately, that his argument is somewhat different to the one given above really becomes clear only in the last section of chapter 7. Consider the sorts of cases where ascribers, with seeming propriety, apply to their subjects the epistemic standards appropriate to those subjects' situations. Such cases are thought to be problematic for contextualism. DeRose argues that this is not so. Say that we've got an ascriber A who is talking about whether a subject *S* knows. In some conversational situations, A's conversational purposes will lead him or her to apply the standards appropriate to his or her situation, for example, when he or she is treating that subject as a source of information. But in other situations, A's conversational purposes will lead him or her to apply the standards appropriate in *S*'s situation, for example, when he or she is discussing whether *S* should take a certain course of action ('*S* should do A because he or she knows that *p*'). Contextualism, or at least DeRose's version, is just the view that what standards an ascriber selects are determined by his or her conversational purposes. So contextualism can easily accommodate both sorts of cases. As DeRose emphasizes, reflection upon our epistemic practice tells us that what epistemic standards we select depends upon our conversational purposes. This is entirely consistent with contextualism, but the defender of SSI needs to insist that, when our conversational purposes demand that we select epistemic standards appropriate to our practical situation, this is a mistake. Our epistemic practice shows that there is a need for what DeRose calls the “flexibility” that contextualism provides.

We think that the above argument for the superiority of contextualism over SSI as an account of our everyday epistemic practice is a good one. However, this argument is certainly insufficient to establish what we earlier identified as the main thesis of TCC. To establish that contextualism provides the best *overall* account of our everyday epistemic practice would require more engagement with the various classical invariantist accounts available and an in-depth analysis of the relative strengths and weaknesses of SSI. DeRose succeeds in presenting a case for contextualism, but it is certainly not a conclusive case.

These critical remarks aside, this is a major and long-awaited work, one that will not disappoint. Even if one has read the papers that are reprinted here, there is still much to be gained by reading them again within the context provided by this book, safe in the knowledge that they are in a form that represents DeRose’s current views. Moreover, the “introductory” chapter 1 is in fact a significant piece of work in its own right, offering a fairly comprehensive history of contextualism along with a fantastic overview of the relevant philosophical terrain and the position of contextualism within that terrain. Anyone serious about epistemology should read this introduction. Given that anyone serious about epistemology should also have already read many of the papers that make up this volume, what you have in this book is thus essential reading for epistemologists *simpliciter*. If there is a problem with this work, it is that once you get to the end, your appetite to start on volume 2 is pretty strong. So let us hope that DeRose doesn’t take too long to present us with the second installment.

Robin McKenna and Duncan Pritchard
University of Edinburgh

Philosophical Review, Vol. 120, No. 3, 2011
DOI 10.1215/00318108-1263719

Elisabeth A. Lloyd, *Science, Politics, and Evolution*.
Cambridge: Cambridge University Press, 2008. vii + 301 pp.

In this remarkable collection of wide-ranging essays by Elisabeth Lloyd, a leading philosopher of biology, we finally have access to a single source that amply exhibits the extent and depth of her fine scholarship. The twelve essays, spanning over twenty-five years, fall naturally into three areas: semantic theory of science, philosophy of (evolutionary) biology, and feminist philosophy of science. Lloyd deftly deals with abstract and technical material in a lucid manner without compromising in any way the complexity of the content, making the book a highly rewarding read for both the specialist and the general scholar. The absence of an introduction may irk some, for others, though, it should make exploring the volume akin to undertaking an expedition without a map, but, fortunately, one that turns out to be a very satisfying adventure indeed.

Up until the late 1970s, the dominant framework for understanding scientific theories was the covering law model of science, or the syntactic view championed by the logical positivists and empiricists. Its dominance was challenged in the 1970s and the early 1980s by the semantic view of science, which sought to give prominence to the role of (mathematical) models in scientific inquiry. In the first three essays, written between 1983 and 1987, Lloyd, an early proponent of the semantic view, aims to highlight its advantages over not only the covering law model but also the Popperian and realist views. The failure of the covering law model to adequately account for evolutionary theory presented a major drawback for its adherents, while bolstering the semantic view, which provides a fine framework for both physical and biological theories. Lloyd's first essay demonstrates the paucity of the model covering law in addition to arguing that views that take Darwin's support for his theory of natural selection to be an instance of inference to the best explanation are misguided largely because an examination of the *Origins* and Darwin's correspondence with his contemporaries shows that it is "the structure of the theory and its relation to empirical evidence," not its explanatory power that "served as Darwin's primary support for his theory" (4).

The second essay examines population genetics, showing the ability of the semantic view to "provide an analytical framework sensitive to the relevant theoretical problems" (21). In the third essay, hypothesis testing and confirmation in evolutionary and ecological theory are explored to argue for "greater complexity and variety" in the semantic approach than in the rival views such as the one due to Popper. Much more on these themes and issues can be found in

I am grateful to Michael Dietrich for comments on an earlier draft.

Published by Duke University Press

Lloyd's book, *The Structure and Confirmation of Evolutionary Theory*, first published in 1988 and later reprinted in 1994.

Lloyd has worked closely with eminent biologists such as Stephen J. Gould and Richard Lewontin, marking a symmetric and fruitful exchange of ideas between biologists and philosophers. The next four essays engage evolutionary biologists and philosophers of biology alike in examining several puzzles and controversies regarding the units and levels of selection (henceforth ULS). Since the time of Darwin, biologists have pondered the question of at what level natural selection occurs: is it at the level of the individual organism, the genes, groups, species, or some combination of these? In "Units and Levels of Selection," which is a classic, four distinct questions are delineated that can aid philosophers and biologists in seeing why it is that the ULS debates seem at times to go nowhere. These questions are the replicator question (issue of the entity which is copied in successive replications), the interactor question (issue of the entity on "which selection acts directly"), the beneficiary question (issue of the entity that "ultimately benefits from the selection process"), and the manifestor-of-adaptation question (issue of the entity that can be said to acquire "adaptation as a result of the selection process"). In addition, Lloyd points to an ambiguity between two uses of adaptation: the "strong" or "engineering" adaptation and the "weak" or "product-of-selection adaptation." With this taxonomy in hand, Lloyd carefully analyzes a variety of central issues in the units of selection controversies to conclude that many of the puzzles that have arisen in this context can be traced to how "numerous players were mixing and matching the various questions in developing their requirements for what it takes to be a real unit of selection" (110). Ultimately, Lloyd does not favor one level of selection over others (although she has long argued for group-level selection) but succeeds in showing how best to disambiguate (with the aid of her taxonomy) some of the heated disagreements and debates surrounding the units and levels of selection. This essay is essential reading for anyone interested in the ULS debates.

In "Species Selection and Variability," coauthored with Stephen J. Gould, the notion of emergent fitness under the "interactor" approach is employed to argue that variability (and other aggregate traits) can be involved in species selection. The role of species variability in "long-term species survival" was recognized by Dobzhansky and others, but in the 1950s evolutionists abandoned this idea, largely because of their failure to distinguish between adaptation and selection processes. That is, they failed to see that all selection processes do not necessarily bring about adaptation. On the basis of their "interactor" approach, with its focus on the relationship between character and fitness instead of just on character (as is the case under the "emergent character" approach), Gould and Lloyd suggest that species-level variability like other group-level properties has been "discredited through its association with unsupportable arguments for group-level adaptations" (89). This paper is another classic.

In "An Open Letter to Elliott Sober and David Sloan Wilson, Regarding Their Book, *Unto Others*: *The Evolution and Psychology of Unselfish Behavior*," the only previously unpublished paper in the volume, Lloyd discusses the issue of pluralism (or the multilevel view of natural selection), according to which any given selection process can be described at many different levels in the biological hierarchy such as genes, organisms, groups, and so on. During the 1960s, the hypothesis of group selection was highly criticized. In their 1998 book, Sober and Wilson argue for the importance of group selection as a cause of evolution, invoking the concept of group selection to explain the existence of altruism in evolution, especially human evolution. Those highly skeptical of group selection as an evolutionary cause have criticized Sober and Wilson for not providing credible evidence. Lloyd appeals to her distinction between "engineering" adaptation and "product-of-selection" adaptation to show exactly why Sober and Wilson's opponents may well be justified. Because much of the opposition to Sober and Wilson's view targets their reliance on the "product-of-selection" notion of adaptation, Lloyd argues, they must either take seriously the notion of engineering adaptation or provide a more convincing argument against it if their view is to enjoy wider acceptance. So far, however, critics of the engineering view have not been outnumbered by its supporters.

The essay "Problems of Pluralism" provides a sharp criticism of "genic pluralism" to defend the legitimacy of the units of selection debates. Unlike Lloyd and Sober and Wilson, who take selection at a given level (say, that of the group) to be an empirical fact, Sterelny, Kitcher, and Waters (SKW), who call themselves the "genic pluralists," argue that there is no such fact and that to demand what the real unit of selection is is to engage in "muddled metaphysics" (109). Essentially, their argument is that any case of selection can be modeled in various different, equally correct ways. Whereas the "genic selectionists" claim that selection can be adequately represented at the genic level, the "genic pluralists" argue that there are multiple, equally adequate models of selection, but that "there is a causally adequate, general evolutionary theory purely at the genic level, one that does not require any appeal to higher-level causal interactors" (110). By showing that SKW's pluralism is "peculiarly weak" and that "genic models are explicitly derived from causal models involving higher-level interactors" (111), Lloyd makes a case for a stronger form of pluralism exemplified in the hierarchical selection model she advocates.

The next two essays examine the problems that beset the careless combination of insights from different fields in biology. In "Normality and Variation: The Human Genome Project and the Ideal Human Type," Lloyd cautions against the sloppy use of insights from molecular biology in medicine. Advocates of the Human Genome Project claim that it will provide the "recipe" to conquer all diseases by virtue of making it possible to define everything at the molecular level. Opposing this view, Lloyd argues that by itself molecular- or DNA-level descriptions are not very useful because what counts as abnormal at the genetic level may not do so at the level of the organism, and vice versa. She correctly challenges the underlying belief in the "giveness" of the categories of normality, health, and disease by reminding us of the social (read "normative") nature of such categories. The next essay, "Evolutionary Psychology," focuses on evolutionary psychology, successor to the now-debunked field of sociobiology, which employs evolutionary theory to explain human psychological behavior. More specifically, Lloyd targets the work of the psychologists Leda Cosmides and John Toobey, whose work on the content effects of the Wason Selection task made a big splash in the late 1980s, arguing against their claim to have provided for their view "empirical evidence that would support a claim of evolutionary adaptation" (168). Lloyd is not against an evolutionary approach to psychology; rather, her position is of a piece with that of Gould and Lewontin in their justly famous "Spandrels" paper, where they rightly object to the "potentially distorting manner in which . . . [adaptive] hypotheses were pursued to the exclusion of other potential explanations" (164).

The last three essays nicely exhibit both Lloyd's feminist take on the politics of science as well as her strong objection to the largely hostile reception given to feminist epistemology and scientific research by both philosophers and scientists. This section should be of special interest to both feminist philosophers and their critics. In particular, "The Double Standard for Feminist Epistemologies" is of great salience for anyone interested in the issue of objectivity or in feminist philosophy of science. The purpose of this excellent essay is to dig deep into the issues surrounding the complex notion of objectivity, with the explicit purpose of demonstrating why analytic philosophers must recognize the crucial role of sex and gender in their work. Lloyd identifies four distinct meanings of "objectivity" and "objective" in the analytic philosophy literature. Thus,

- Objective can mean detached, disinterested, unbiased, impersonal, invested in no particular point of view . . .
- Objective means public, available, observable, or accessible . . .
- Objective means existing independently or separately from us;
- Objective means really existing, Really Real, the way things really are. (173)

While the above meanings are distinct, all appear in a certain account of objectivity that Lloyd labels the "ontological tyranny." This is the "strong claim that 'objective' reality—the reality converged on through the application of objective methods—equals all of the Really Real" (177). Put differently, the notion of objectivity as disinterested rests on the "ontological tyranny" that it is possible for the knower to be completely detached from the reality he or she seeks to know. The origin of this position can be traced back to the seventeenth-century distinction between primary and secondary qualities. In the eighteenth century, it turns up in the guise of the "philosophical folk view of objective knowledge"

(179). In its contemporary version, it is discernible in the widespread view Lloyd labels "Type/Law Convergent Realism," which assumes that "real objectivity will result in a convergence on One True Description" of reality (181). Most contemporary philosophers deny the possibility of knowing a reality that exists independently of us. Among these, Lloyd examines the views of objectivity and knowledge expounded by Pierce, Williams, Nagel, McDowell, Carnap, and Searle, who in one sense or another reject this ontological tyranny, thereby making room in their accounts for "the potential relevance of a variety of social considerations" (188). Through a close reading of their work, Lloyd shows that all these thinkers accept, to a certain degree, the legitimacy of social factors in matters concerning objectivity and objective knowledge. Thus, by their own standards, they ought to allow sex and gender as legitimate units of analysis in their accounts of objectivity; however, none do. Moreover, holding feminist scholarship, with its emphasis on the crucial role of values in science, to the standard of the "ontological tyranny" that they themselves reject, argues Lloyd, constitutes a "double standard." This essay, published in 1995, should prove to be a great resource for more recent debates and further work on scientific objectivity.

The late 1980s and early 1990s saw the rise of what came to be known as the "science wars," in which scientists and social scientists accused scholars of social studies of science (both feminist and antifeminist) of being antiscience. To set the record straight and continuing her argument from the previous essay, in "Science and Anti-Science" Lloyd carefully examines the strategies (and rants) of these critics aimed at "discrediting or excluding" scholars of social studies of science. By gleaning impressive textual evidence in support of her argument that science studies scholars abide by the standards of evidence and argumentation that meet all the requirements deemed necessary by science itself, Lloyd concludes that these scholars are wrongfully charged with being antiscience when it is, in fact, the critics themselves who "stand as the true enemies of rationality, objectivity, and the long-term success of the sciences" (254). This is primarily because the critics fail to see that social accounts of science are not meant to show that only the social (and political) aspects of science are relevant. Rather, the argument is that any account of science that ignores its social and political dimensions is incomplete. Moreover, focusing on these hitherto-neglected aspects of science is invaluable in the "self-correction" of science itself.

While this essay is critical in making sense of the exclusion and marginalization of social accounts of science, what remains puzzling, however, is that in her quest to align feminist accounts of science with those of nonfeminist scholars (mostly historians and sociologists of science), Lloyd ignores the crucial distinction as well as the tension between feminist and nonfeminist scholars of the social studies of science. This distinction is evident in the insistence of the nonfeminist science studies scholars in asserting that their overall project is essentially descriptive in character, whereas feminist scholars recognize (in fact, even insist on) the importance of taking a normative stance. This may in part explain the unexplained absence of sex and gender analysis from the nonfeminist social studies of science project itself, which otherwise embraces all and any social factors in its study of science. Here one can discern a symmetry (in their attitude to sex and gender as analytical categories) between the approach of the analytic philosophers criticized in the previous essay and the nonfeminist social studies of science scholars that Lloyd welcomes as allies of feminists. It is at this juncture in the volume that it would have been helpful if Lloyd had updated her stance. Furthermore, there are several kinds of questions that would be appropriate to raise here. Do sex and gender inevitably influence scientists' knowledge discourse? Is the role of sex and gender in scientific accounts (even if not inevitable) always undesirable? Does finding sex and gender bias (aka androcentrism) in science exhaust feminist scholarship? Recent feminist scholarship takes us beyond the earlier feminist critical stance. Lloyd's input here would have been most useful.

The last essay, "Pre-theoretical Assumptions in Evolutionary Explanations of Female Sexuality," is a precursor of the longer argument presented in *The Case of the Female Orgasm*.[1] It provides a highly engaging test case for advancing the arguments in the previous two essays, as well as provides the argument, discussed earlier, against the uncritical use of adaptationist hypotheses. Lloyd's aim is to examine evolutionary arguments to see what they say about *whether* and *how* women's sexuality is related to reproduction (255). Of the thirteen studies by sexologists that Lloyd examines, all but one find a strong connection between female sexuality and reproduction. This connection, however, comes apart when the studies are held against the standards of empirical adequacy and consistency. More specifically, these studies attempt to establish that female orgasm is an adaptive trait. However, the only study that gets her approval, conducted by Donald Symons, concludes that female orgasm is a by-product of the selection for male orgasm. That male orgasm is an adaptation is uncontroversial, but, as evidence shows, to argue that female orgasm could not be otherwise is misguided. Lloyd argues, and I think rightly, that underlying the studies that claim that female orgasm is an adaptation can be found the adaptationist bias of the sexologists in conjunction with the androcentric bias that ties women's sexuality to their reproductive function. Those unconvinced by this argument are well advised to examine Lloyd's meticulous arguments in her above-mentioned book on the subject. Interestingly, Lloyd has come under attack by several feminists who take her study to endorse the view that female orgasm was *just* a by-product (a *mere* derivative) of the male orgasm.

1. Elisabeth A. Lloyd, *The Case of the Female Orgasm: Bias in the Science of Evolution* (Cambridge, MA: Harvard University Press, 2005).

Overall, this book is of immense value for many reasons, not least of which is that reading Lloyd's work on varied topics in a single volume lends coherence to her different interests. Her highly valuable contributions to philosophy of biology, long recognized by specialists, can now be appreciated by generalists as well. But most of all, Lloyd's contribution to feminist philosophy sets this volume apart from other works in philosophy of biology. It should prove highly useful for a graduate course in philosophy of biology or philosophy of science that refuses to be too narrow.

Neelam Sethi
Cornell University

Philosophical Review, Vol. 120, No. 3, 2011
DOI 10.1215/00318108-1263728