

- Rapaport, William J., & Shapiro, Stuart C. (1984), "Quasi-Indexical Reference in Propositional Semantic Newtorks," *Proc. 10th Int'l. Conf. Computational Linguistics (COLING-84; Stanford University)* (Morristown, NJ: Assoc. for Computational Linguistics): 65-70.
- Sayre, Kenneth M. (1986), "Intentionality and Information Processing: An Alternative Model for Cognitive Science," *Behavioral and Brain Sciences* 9: 121-165.
- Searle, John R. (1980), "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3: 417-57.
- Searle, John R. (1983), *Intentionality: An Essay in the Philosophy of Mind* (Cambridge: Cambridge University Press).
- Searle, John R. (1984a), "1984 Reith Lectures—Minds, Brains and Science: 1. A froth on reality," *The Listener* 112 (8 November 1984): 14-16.
- Searle, John R. (1984b), "1984 Reith Lectures—Minds, Brains and Science: 2. Beer cans and meat machines," *The Listener* 112 (15 November 1984): 14-16.
- Searle, John R. (1984c), "1984 Reith Lectures—Minds Brains and Science: 3. Grandmother knew best," *The Listener* 112 (22 November 1984): 16-18.
- Searle, John R. (1984d), "1984 Reith Lectures: 4. A walk to Patagonia," *The Listener* 112 (29 November 1984): 8, 10-11.
- Searle, John R. (1984e), "The Reith Lectures: 5. A changing reality," *The Listener* 112 (6 December 1984): 15-18.
- Searle, John R. (1984f), "1984 Reith Lectures: 6. The freedom of the will," *The Listener* 112 (13 December 1984): 10-12.
- Shapiro, Stuart C., & Rapaport, William J. (1986), "SNePS Considered as a Fully Intensional Propositional Semantic Newtork" *Proc. 5th National Conference on Artificial Intelligence (AAAI-86; Philadelphia)*, Vol. 1 (Los Altos, CA: Morgan Kaufmann): 278-83.
- Shapiro, Stuart C., & Rapaport, William J. (1985). "SNePS Considered as a Fully Intensional Propositional Semantic Network," in G. McCalla & N. Cercone (eds.), *The Knowledge Frontier* (New York: Springer-Verlag.): 262-315.
- Tenenbaum, Aaron M., & Augenstein, Moshe J. (1981), *Data Structures Using Pascal* (Englewood Cliffs, NJ: Prentice-Hall).
- Watts, Geoff (moderator) (1984), "The Reith Lecturers 1984 Discussion: Are minds really free?," *The Listener* 112 (20-27 December 1984): 46-47.

Daniel C. Dennett, *Elbow Room: the varieties of free will worth wanting* (Cambridge, Mass.: The M.I.T. Press, 1984), x + 200 pp., \$8.95.

PETER VAN INWAGEN

SYRACUSE UNIVERSITY

The announced purpose of this book is to dispel the confusions and anxieties by means of which philosophers persuade one another that there is such a thing as "the problem of free will." For that is what they do: If they

didn't, everyone would long ago have turned his attention to more profitable areas of philosophical inquiry. Despite the fact that a "compatibilist" treatment of the so-called problem is obviously correct (Dennett maintains), philosophers are still arguing about it; the arguments, therefore, are evidently a mere facade. Behind the facade lurk fears—the fear of naturalism; the fear that one is a mere automaton; the fear that one is not really a *self*—fears begotten by conceptual confusion on the desire to have something to write articles about.¹ These misbegotten fears are not viable, but they are kept alive by extraordinary measures: continual transfusions of plausibility from really fearsome things by misapplied "intuition pumps." My impression is that Dennett regards even the traditional compatibilist as prey to the fears that collectively make up the free-will problem: if the compatibilist weren't prey to these fears, he wouldn't bother arguing with the incompatibilist. (Dennett does not say much about other compatibilists, and he certainly does not acknowledge the predominance of the doctrine among contemporary writers on free will. Why Dennett does not tell his readers that the *content* of his views on free will—as opposed to his views on the etiology of the free-will debate—is the prevailing orthodoxy, I'm not quite sure. I *am* sure that a reader of *Elbow Room* who was not familiar with the history of philosophy in the English-speaking countries would not learn from it that most analytical philosophers are compatibilists. Because Dennett says so little about his fellow compatibilists, my description of his attitude toward them is largely guesswork. My impression is that he regards them with the sort of condescension that Marxist and Freudian atheists reserve for old-fashioned Enlightenment atheists like d'Holbach and Russell: such people don't understand the real sources of the superstition they fulminate against and are therefore at best imperfectly free of it themselves.) Like many cultural diagnosticians, Dennett sees symptoms everywhere. (I am put in mind of a Freudian analysis of science-fiction I once read, in which a spaceship was a phallus if it was longer than it was wide, and a womb otherwise.) An incompatibilist can't so much as mention a man locked in a room without Dennett's accusing him of insidiously suggesting that if determinism were true it would be just as if we were all in jail.² Dennett also has an annoying trick of asking rhetorical questions of the *What do you suppose it is that makes him beat his wife so much?* type. What is it, he wonders rhetorically, that people find so horrible about determinism? (See, e.g., p. 15. But the idea that philosophical rejections of determinism are based on fear is one of the central ideas of the book, and it surfaces repeatedly.) Having raised the question, he proceeds to answer it. He does not raise the question whether anyone actually does find determinism horrible (as opposed to false, ill-supported, or in conflict with some thesis that seems antecedently more probable than determinism). As long as we are speculating about people's hidden motives, I will speculate that philosophers who speculate about the hidden motives of other philosophers do so because it is easier to speculate about the invisible and unverifiable than to address arguments. (And, of course, it is very pleasant to represent oneself as someone who is in a position to expose the hidden motives of others. I'm certainly en-

joying it.) Is this fair to Dennett? Let's say that it is as fair as speculation about hidden motives ever is.

Despite all this diagnostic nonsense—I think it's nonsense, anyway—this is a rather good book. A lot of confusions that some people are doubtless prey to are nicely straightened out³ (I particularly commend the beautiful exposure in Chapter Three of certain widespread confusions about control and self-control), and a wealth of interesting empirical information finds its way into the examples. But I don't think the book does much to advance our understanding of the traditional problem of free will. In the sequel, I shall explain my reservations.

Dennett's discussion of what I would call "the problem of free will and determinism" occurs in Chapter Six, the chapter called "Could Have Done Otherwise." I have had a *very* hard time deciding what the argument of this chapter is. I will present my own rational reconstruction of parts of the argument. Doubtless it is no more accurate than most rational reconstructions.

As I read Dennett, Chapter Six is mainly an attack on the principle
 CU An agent is morally responsible for an act only if that act was causally undetermined.

But this is a controversial interpretation. Dennett's *announced* target is stated in more or less the following words (he gives no "official" statement of it):

CDO An agent is morally responsible for an act only if that agent could have done otherwise (*sc.* than perform that act).

What Dennett *means* by 'X did A and could have done otherwise', however, seems to be something like 'X's doing A was causally undetermined'. At any rate, I shall argue that this is what he has to mean by this phrase if his argument against CDO is even to be relevant to its conclusion (much less *valid*). It might be objected that I must be wrong about this, since, according to Dennett, he proposes to attack a "widely accepted" principle—and CU, far from being widely accepted, is an object of popular derision. I answer that Dennett *thinks* he is attacking a widely accepted principle and he is wrong.

Dennett's mistake is a complicated one. In broad outline, it is this. Consider the principle

WA An agent is morally responsible for an act only if that agent was able to do otherwise.

The principle CDO is ambiguous: it might mean either CU or WA. WA is widely accepted. But it is to CU that Dennett's argument applies.

Dennett's argument, or an important part of it, is something like this.

When philosophers employ the principle CDO, they mean that the agent in question is morally responsible only if he could have done otherwise *in exactly the same circumstances*. But it is in practice impossible to find out whether someone could have done otherwise than he did in *exactly* the same circumstances, for "circumstances" in this context refers (in defiance of its etymology) not only to *external* circumstances but also to

the state of the agent's body and brain. If CDO were correct, therefore, we could never know whether an agent was responsible for anything. If a concept is such that we can never, in practice if not in principle, know whether it applies, then that concept, however interesting it may be to philosophers, is not a concept that would be of any use in everyday life. Therefore, if there is *any* concept that goes by the name "moral responsibility" and which satisfies the demands of CDO, it is not the concept that goes by that name in everyday life.

I shall presently criticize this argument. But first I will show that it must be read as an attack on the disreputable principle CU and not on the respectable principle WA.

It is well known to students of the free-will problem that 'could have' is ambiguous: these words can mean either 'might have' or 'was able to'. This ambiguity can be especially acute when 'could have' occurs in the consequent of a conditional. Consider the sentences, 'If you'd startled him, he could have fallen' and 'If you'd warned him, he could have escaped'. These sentences mean, respectively,

You startled him $\diamond \rightarrow$ he fell

You warned him $\square \rightarrow$ he was able to escape.

In the former sentence, 'could have' is absorbed into the subjunctive connective; in the latter, into the indicative consequent. The following two cases (adapted from Austin) provide a second example of this ambiguity. In each case, one financier addresses another. *Case One* (could have = were able to): "You could have ruined me this morning. I want you to know that I'm grateful." *Case Two* (could have = might have): "You could have ruined me this morning. Warn me the next time you're going to pull a stunt like that." It is important to remember that the proposition that X was able to do otherwise does not, uncontroversially, entail the proposition that X might have done otherwise in exactly the same circumstances. Those who say that this entailment holds are called *incompatibilists*. Those who deny that it holds are, of course, called *compatibilists*.

Now recall Dennett's argument. According to Dennett, what the proponents of CDO really mean by their thesis is more completely expressed in these words:

ESC An agent is morally responsible for an act only if that agent could have done otherwise in exactly the same circumstances.

Consider the consequent of this conditional (call the agent 'X'):

X could have done otherwise in exactly the same circumstances.

Do the words 'could have done' in this sentence mean 'was able to do' or 'might have done'? They must mean 'might have done'; they cannot mean 'was able to'. Or, at any rate, charity forbids the latter interpretation. For suppose 'could have done' did mean 'was able to'. Then this sentence would read

X was in exactly the circumstances X was actually in $\square \rightarrow$ X was able to do otherwise.

But since the antecedent of this conditional is “automatically” true, what is the point of it? What is the point of saying, e.g., ‘If the election had turned out exactly the way it did, Reagan would have been elected’? Why not simply say, ‘Reagan was elected’? What is the point of saying ‘If he had been at t exactly as he then was, he would (at t) have been able to do something other than what he is fact proceeded to do’? Why not simply say, ‘He was (at t) able to do something other than what he in fact proceeded to do’?⁴ (But it makes perfect sense to say, ‘(Even) if he had been at t exactly as he then was, he might have done something other than what he in fact proceeded to do’.) Charity dictates that we assume that people do not add pointless qualifications to what they say—especially when they insist on the importance of those qualifications. Charity dictates, therefore, that we take ESC to mean

An agent is morally responsible for an act only if that agent might have done otherwise in exactly the same circumstances.

To say that an agent might have done otherwise in exactly the same circumstances is to say that if the agent and his environment were returned to *exactly* the state they were in just before he acted, then he might well act differently the “second time round.” And to assert that conditional is to say that what the agent did in actuality was causally undetermined. For if what he did in actuality was causally *determined* by his inner state and the state of his environment at t , then returning him and his environment to *exactly* their condition at t would have to produce the *same* result the second time round.

Therefore, the real target of Dennett’s argument is not the respectable principle WA (or any other widely accepted principle), but CU, a principle most philosophers reject. Is his argument against this unpopular principle cogent? In order the better to answer this question, let us recast it as an argument that is explicitly directed against CU:

Since we can never in practice return an agent’s body, brain, and environment to exactly the conditions they were in at some previous time—since, in fact, we cannot even trace in any detail the workings of the brain—, it is in practice impossible to find out whether an agent’s act was causally undetermined. If CU were correct, therefore, we could never know whether an agent was responsible for anything. If a concept is such that we can never, in practice if not in principle, know whether it applies, then that concept, however interesting it may be to philosophers, is not a concept that would be of any use in everyday life. Therefore, if there is *any* concept that does by the name “moral responsibility” and which satisfies the demands of CU, it is not the concept that goes by that name in everyday life.

(We may note that anyone who accepts this argument should reject the contention of Broad, Hobart, Ayer, Nowell-Smith, Smart, and many others, that an agent can be morally responsible for an act only if that act was determined.⁵ For it is exactly as hard to find out whether a given

act is determined as it is to find out whether that act is undetermined.)

This argument seems to me to be invalid. Take any concept that we are perfectly sure we do employ in everyday life. The concept of a dog, say. That there are dogs entails the truth of Goldbach's Conjecture, or else entails the falsity of Goldbach's Conjecture. It is, in practice if not in principle, impossible for me (and you, too) to find out whether Goldbach's Conjecture is true or false. But it hardly follows from these facts that the concept of a dog is of no interest to the working veterinarian or burglar.

Now it might be held that this reply works by exploiting the "paradoxes of strict implication" or some other defect in the orthodox philosophical understanding of entailment. This rejoinder, however, is directed at a merely accidental feature of the above criticism of Dennett's argument. Consider "rogs". Rogs are robotic "dogs"—creatures of the philosophically invaluable Martians—, which we cannot (in practice) tell from standard, protein dogs. Let DR be the principle, "A thing is a dog only if it is not a rog." We argue:

It is in practice impossible to find out whether a (superficially doggish) thing is a rog or a dog. If DR were correct, therefore, we could never know whether a (doggish) thing was a dog . . . Therefore, if there is *any* concept that goes by the name "dog" and which satisfies the demands of DR, it is not the concept that goes by that name in everyday life.

The conclusion of this argument is false, since a thing that looked like a dog but which was really a clever robot would *not* be a dog. The argument is therefore defective. What is wrong with it? It is important to keep in mind that this is not a *skeptical* argument. Its conclusion is not the (false) proposition that we never know whether anything is a dog; it is, rather, the (false) proposition that the everyday concept "dog" is compatible with "robot." Nevertheless, this argument shares certain premises with the standard "wild hypothesis" arguments for skepticism. It shares, for example, the premise that if p entails q , then one cannot come to know that p unless one first (or, at least, simultaneously) finds out whether q . This principle has come in for a good deal of deserved abuse. Dennett's argument against CU—or CDO or whatever its target is—would appear also to have this principle as a premise. I would suggest that *whatever* the defect in the rog/dog argument is, Dennett's argument shares that defect.

It may be suggested that the two arguments are not really parallel, owing to the fact that, while the thesis that our acts are causally determined is not absurd, not something we are justified in assigning an infinitesimal subjective probability to (some even find it highly plausible), the thesis that there are rogs at large in our environment is absurd. I accept these two judgments. But I would point out that neither the thesis that human acts are causally determined nor its denial had occurred to anyone during the formative years of the concept of moral responsibility. Therefore, *no* fact about this thesis (whether it be the fact that it does not seem to 20th-century, middle-class, naturalistically inclined professors of philosophy to be absurd, or any other fact about it) had any opportunity to play a causal

role in the development of the concept of moral responsibility. If, in the future, some amazing revelation transpired that made us think that there might well be rogs abroad in the land, this prodigy would not change the fact that the rog/dog argument is a bad argument for a false conclusion.

Dennett has a second argument against CU—or at least it is most charitably interpreted as an argument against CU. He imagines a robot that has performed badly in some unanticipated way in a certain situation. He imagines its designers thereupon asking one another, “Could it have done otherwise?” (He *says* that they use those words. It’s his case; but it doesn’t seem to me that people in those circumstances would be likely to say that.) What the designers would be asking, Dennett suggests, is whether the robot would behave in the same undesirable way in circumstances that were *for all practical purposes* the same: if it would have behaved in the same undesirable way in relevantly similar circumstances, then it “couldn’t have done otherwise.” (The designers would grant—if they thought about it—that the robot would behave in the same way in *exactly* the same circumstances; but since the chance of an *exact* duplication of the robot’s original situation is infinitesimal, it will probably not occur to them even to formulate this thesis.) If the robot “couldn’t have done otherwise,” the designers have a good reason to modify or reprogram the robot or its successors. Dennett’s point, apparently, is that this case is analogous to cases in which we ask the question ‘Could he have done otherwise?’ about a human being, and he strongly suggests that what we really want to know when we ask whether a human being could have done otherwise is this: Should we modify him (or his successors) in such a way that in relevantly similar circumstances he (or they) will do better?—or, if we are not in a position to effect such modifications, should he be kept out of relevantly similar circumstances in the future? We are *not* raising the question whether we should modify him to deal with (or should keep him out of) a recurrence of *identical* circumstances, since the probability of a recurrence of identical circumstances is effectively zero. Therefore, (finally), the answer, whatever it may be, to the philosophical question, ‘Would X have done otherwise in exactly the same circumstances?’, has no bearing on the everyday question ‘Could X have done otherwise?’

Sometimes something like Dennett’s account of the everyday use of ‘Could he have done otherwise?’ may be right. We send a spy into Ruritania, having given him certain training and instructions, and he is immediately captured. “Could he have done better?” we ask, meaning, Should the training and instructions of his successors be different? And we might indeed use these words, though it seems to me to be far more probable that we’d say something along the lines of, “Was that sheer bad luck, or do we need to review our training procedures?” But remember that our general topic is the responsibility of the *agent* (no pun intended). And if we ask, concerning the spy, “Could he have done better?” in Dennett’s sense, we are not even raising the question whether *he* was responsible for the failure to his mission—though we may proceed to raise the question whether the Director of Espionage Training was. Contrast this use of ‘could have done better’ with its use in the following exchange.

M says, "I'm inclined not to blame our preparations, but 007 himself. It's his fault. He could have done better." A friend of 007's defends him: "Really, M. The pilot dropped him in the middle of a KGB battalion. There was nothing he could have done." The friend is saying that, since 007 did not have it within his power to do better, he cannot be blamed.⁶ And this returns us to the question of incompatibilism: Does the everyday use of 'could have done [= was able to do] otherwise' entail 'Might have done otherwise in exactly the same circumstances'? There are arguments for incompatibilism. Perhaps they are defective, but, if so, the reader will not find out about their defects (or their content or their existence) from Dennett's book.

I have tried to show that Dennett's arguments against CDO must, since they contain the qualification "in exactly the same circumstances," be understood as arguments against CU and not against WA; and that, considered as arguments against CU, they are unconvincing.

What about WA? Has Dennett any arguments that show, or which could be modified to show, that WA is false? It is clear that he thinks he has such an argument. As I have said, he thinks that his target is widely accepted; and at one point (p. 132) he identifies his target with the principle that Harry Frankfurt has attacked under the name "the Principle of Alternate Possibilities."⁷ Dennett's belief that, in the arguments discussed above, he is attacking a widely accepted principle is a result of his failure to make appropriate distinctions among the similar-sounding principles WA, CDO, ESC, and CU. But Frankfurt's "Principle of Alternate Possibilities" is just WA. Since Dennett gives an accurate transcription of Frankfurt's argument against WA, and—while not disputing the cogency of Frankfurt's argument—promises a more radical attack on the same principle, I think he does mean to attack WA "under that description."

Dennett has only one argument that does not involve the fatal qualification 'in exactly the same circumstances' and which is, therefore, not automatically inapplicable to WA. Dennett asks us to consider Martin Luther, who—according to popular, if not to academic, history—said, "Here I stand. I can do no other." Several people have told me that they find Dennett's interpretation of Luther's words puzzling. Surely, they protest, what Luther meant (or what someone in Luther's position would mean if he in fact said those words) is this: This is the only morally permissible course of action open to me. Surely—their protest continues—these words cannot be taken to mean, 'I am, literally, *unable* to do otherwise'. It is certainly true that when someone says "I can't," he often means "I can't without . . .," where the content of the ellipsis is supplied by context. Luther (might have) meant "I can't do it without acting wrongly." The bank officer who tells you that he can't approve your request for a loan does not mean that it isn't within his power so to arrange matters that your request is approved; he means that it isn't within his power to do this without violating the bank's rules (and thereby putting his job in jeopardy). I think that there is probably something to this reaction to Dennett's reading of Luther's words, but it would be a mistake to press the point, since other cases could be found that would serve Dennett's

purposes. Cato the Elder and George Washington are both supposed to have been unable to lie. (Mark Twain: "I am morally superior to George Washington. He couldn't tell a lie. I can and I don't.") Suppose Washington refused to tell a lie on a certain occasion. Don't we admire him, even if we believe he wasn't able to lie? Don't we regard his telling the truth (in the face of a severe spanking for chopping down the cherry tree) as an act for which he was morally responsible? I despair of trying to answer these questions adequately in a few sentences. There are too many things that would have to be discussed first. I will briefly mention a few of them. First, I doubt whether we ever do, in ordinary moral discourse, "hold people responsible for their acts." We hold them responsible for the *results* or *consequences* of their acts (or failures to act). 'Just look at what you've done' means 'Just look at what you've caused.'⁸ Secondly, talk of moral responsibility is much more closely tied to expressions of blame than to expressions of gratitude. (People often talk as if 'praise' were the opposite of 'blame'. It is not: you can praise Einstein's genius, but you can't blame a moron for being one. 'Blame' has no real opposite in English. 'Gratitude' is a better candidate than 'praise', but it is not satisfactory, since, although there is impersonal, disinterested blame, there is no such thing as disinterested gratitude.⁹) Thirdly, I believe that the staunchest proponents of WA admit that this principle is loosely stated, and that it should be elaborated to take account of "drunk driver" cases: cases in which a person's present inability to do otherwise is due to past misbehavior. (He was unable to swerve fast enough to avoid hitting the taxi at *t*, but we still hold him responsible for the accident, since there was a time at which he could have so arranged matters that his reflexes were unimpaired at *t*—or that he was not at *t* in a situation that called for unimpaired reflexes.)

One who takes these three points seriously (as I do) may well want to say (as I do) that WA *is* false, but that it nevertheless expresses an important moral insight, one that is better expressed as follows:

No event (or state of affairs) can be X's fault, or even partly X's fault, unless there was a time at which X was able so to arrange matters that that event (or state of affairs) not occur (or obtain).

Now this principle, too, can be disputed on various grounds.¹⁰ (A philosopher is someone who thinks otherwise.) But I think it is a quite reasonable principle and that it seems no less reasonable after I have examined its consequences for the cases of Luther, Cato, and Washington. I think, moreover, that it can do all the philosophical work that anyone has ever expected WA to do.

I conclude that, despite its considerable merits, *Elbow Room* has nothing to tell us about whether the thesis that we are morally accountable for the results of our acts is compatible with the thesis that those acts and their results are logical consequences of the laws of nature and certain propositions about the remote past.

NOTES

¹Or perhaps on a humanistic hatred of science. The immediate ancestry of the fears is not consistently described, but there is certainly bad blood in the family

²I once needed to make the point that someone can *make* a choice without *having* a choice; I remarked that a man who is locked in a room without knowing it can make a choice about leaving without having a choice about leaving. Dennett (p. 105) cites this harmless piece of conceptual analysis as an instance of "the well-worn intuition pump, the Invisible Jail." Dennett's rule, apparently, is that any incompatibilist who uses a story involving a locked room for any conceptual purpose whatever is subliminally suggesting that a deterministic world would be like a prison.

³A recent book, *Free Will and Responsibility* by Jennifer Trusted (Oxford: 1984), might have been written to prove that there are people who are prey to these confusions.

⁴Or put the point like this: to say that someone was (at *t*) able to raise his hand is simply to say that at *t* he possessed a certain property. It makes no more sense to say 'He was (at *t*) able to raise his hand in exactly the same circumstances' than it does to say 'He was (at *t*) fat in exactly the same circumstances'.

⁵And Dennett does reject this contention. See "On Giving Libertarians What They Say They Want," which is Ch. 15 of *Brainstorms* (Bradford Books: 1978).

⁶Here, of course, I am insidiously suggesting that living in a deterministic world would be like being at the mercy of a hostile armed force.

⁷"The Principle of Alternate Possibilities," *The Journal of Philosophy* (1969).

⁸Cf. P. H. Nowell-Smith, "Action and Responsibility," in *Action Theory*, ed. by Myles Brand and Douglas Walton (D. Reidel: 1976) pp. 311-322; cf. esp. p. 315f.

⁹"We give thanks to Thee for Thy great glory," says the *Gloria in Excelsis*. Perhaps this is poetic license.

¹⁰The most troublesome involve "overcausation": X and Y, acting independently, perform acts either of which would be sufficient for the death of Z. . . .

G.P. Baker & P.M.S. Hacker, *Scepticism, Rules & Language* (Oxford: Blackwell 1984), xiv + 140 pp., £ 15.00.

PAOLO LEONARDI

UNIVERSITÀ DI VENEZIA

Baker and Hacker, two among the most well-known Wittgensteinian scholars, have here missed the opportunity to clarify an issue—rules and rule-following—which is both interesting in itself and a central topic in Wittgenstein's *Philosophical Investigations*. From the very title, Baker and Hacker claim to be widening the scope and changing the focus from Wittgenstein, and from the recent debate on Wittgenstein's views on rules and rule-following, in which Kripke took an important and intriguing stand with his *Wittgenstein on rules and private language* (henceforth referred to, in its book version, as *K*).¹