

PHILOSOPHY
AND
PHENOMENOLOGICAL
RESEARCH

A Quarterly Journal

FEATURES OF INTENTIONAL ACTIONS

BY

D. C. DENNETT

REPRINTED FROM VOL. XXIX, NUMBER 2, DECEMBER 1968

Copyright 1968, University of Buffalo

Printed by Wolters-Noordhoff Publishing - Groningen

FEATURES OF INTENTIONAL ACTIONS

I

In recent years there has been a growing chorus claiming that human action, more specifically intentional action, can be so removed from the sphere of physical objects, states, events, and causes that something akin to the "naturalistic fallacy" in ethics can be attributed to anyone who speaks of reasons as causes, or the causes of action, or who worries about the old bugbear of determinism. First soprano in the chorus is Elizabeth Anscombe, and in *Intention* she offers an argument to show that "an action is not called 'intentional' in virtue of any extra feature which exists when it is performed."¹ This argument holds a central place in the support of the new doctrine about action, and I wish to show that the conclusion of the argument, on the most likely interpretation, is both remarkable and untenable. A more modest version of her conclusion is, I believe, tenable, but hardly worth the effort Miss Anscombe expends in providing her proof and incapable of supporting the new doctrine about action.

The argument is presented as a straightforward *reductio ad absurdum* and with most of it I have no quarrel:

That an action is not called 'intentional' in virtue of any extra feature which exists when it is performed, is clear from the following: Let us suppose that there is such a feature, and let us call it '*I*'. Now the intentional character of the action cannot be asserted without giving the description under which it is intentional, since the same action can be intentional under one description and unintentional under another

That is, with just one "feat of motion" I may be both sawing a plank and waking the baby, but only under the former description is my action intentional.

. . . It is however something actually done that is intentional, if there is an intentional action at all. A man no doubt contracts certain muscles in picking up a hammer; but it would generally be false to call his contraction of muscles the intentional act that he performed

¹ Blackwell's, 1957, Second Edition, 1963, pp. 28-29. All citations from Anscombe are from these pages.

Were this not so, it should be possible, for example, to order a man to contract just those muscles, which is generally not possible. The man is not aware, usually, that he is contracting these muscles in just this way, and this disqualifies the contracting as an intentional action. (See Anscombe, pp. 11-13.)

... This does not mean that his contraction of muscles was unintentional. Let us call it 'preintentional'. Are we to say that *I*, which is supposed to be the feature in virtue of which what he does is an intentional action, is something which accompanies a preintentional action, or movement of his body? If so, then the preintentional movement + *I* guarantees that *an* intentional action is performed, but which one? Clearly our symbol '*I*' must be interpreted as a description, or as having an internal relation to a description, of an action....

It is not altogether clear what Anscombe means by an internal relation here, but relying on the dictates of tradition and the context of her argument we can uncover at least a necessary condition for internal relations. A symbol '*I*' would have an internal relation to a description of an action provided it uniquely determined such a description without itself being a description of an action. Thus "the thought: pick up the hammer," or "his volition to pick up the hammer," and perhaps "a sentence about picking up a hammer" would all be expressions that are not descriptions of the action of picking up a hammer, but have an internal relation to such a description.

... But nothing about the man considered by himself in the moment of contraction of his muscles, and nothing in the contraction of the muscles can possibly determine the content of that description which therefore may be *any* one, if we are merely considering what can be determined about the man by himself in the moment.

The conclusion, then, is that since no feature we can detect in the man and his motion could justify our calling it "*I*₁" (e.g., "the volition to saw the plank") rather than "*I*₂" ("the volition to wake the baby"), no such feature could be that in virtue of which we call the action of sawing the plank intentional. This argument leads Miss Anscombe to conclude that "the assumption that some feature of the moment of action constitutes actions as intentional leads us into inextricable confusions, and we must give it up."

The difficulty I see in the argument has to do with the claim that "nothing about the man . . . can possibly determine the content of that description." She offers no support for this claim, and it does not seem to me to be obvious. I suggest that this claim can be interpreted in two different ways, depending on how we interpret her conclusion. and that on one interpretation the claim is insupportable, and on the other in-

sufficient for the argument. First we must look at the different interpretations of her conclusion, that no extra feature is, or can be, the *criterion* for calling an action intentional, and this requires clarification of the notion of a criterion.

II

There is a distinction that is often made between a *symptom* and a *criterion*. Spotting a symptom of *X* gives us evidence for *X*, but the relation between symptom and *X* is held to be contingent; one *could* display the symptom without being *X*. The relation between a criterion for *X* and *X* is, on the other hand, necessary, so that if *A* satisfies the criterion for *X*, that settles the matter: *A* is *X*. Sometimes we decide for various reasons to make what had hitherto been a symptom the criterion for *X*. The doctrine is, I believe, that when we do this, we have *changed concepts*, or "*X*" has suffered a *change in meaning*. This doctrine is plausible when what had been our symptom was only fairly reliable. Thus if jogs on the lie detector polygraph were made the *criterion* for lie-telling, then we would mean something quite different by lie-telling from what we now mean. But where a symptom is strongly reliable (almost always accompanies *X*) or is held to be, as a point of some theory, superreliable (always accompanies *X*), making the symptom the criterion does not seem to change the concept at all. An example that so neatly parallels the case of intentional actions that I shall use it in several ways is that of magnetism. The "old" concept of magnetism had a behavioral criterion — roughly, if you brought iron near *M*, and the iron was drawn to *M*, *M* was a magnet. Then the scientists discovered a feature of molecular alignment (let us call it *M*-alignment) common to all magnets. Did they, as a matter of historical fact, then change concepts, making *M*-alignment the criterion for magnetism? To find out, one would have to ask the scientist what he would call a hypothetical object which had the *M*-alignment but did not attract iron. Would he call it an *M*-aligned nonmagnet, or a magnet that did not work? The question would probably be of only faint interest to the scientist, since he does not expect ever to come across such an object, and if he did, he would have more important things to worry about than what to call it. The scientist is not being annoyingly unphilosophical in taking this attitude, for if his theory — the theory which explains the fact that all iron-attracters are *M*-aligned — is satisfactory, nothing of importance hinges on making the decision about the definition of the term "magnet," which might as well be made by flipping a coin. If we must speak of a change of concept where a superreliable symptom is made the criterion, let us call it a *weak change* (and let us treat "weak" as syncategorematic in this in-

stance, so that a weak change is not a change that is weak, any more than an expectant mother is a mother who is expectant).

If the distinction between symptom and criterion is now clearly drawn, we must still examine what it means to speak of our *actual* criteria for the application of a term, for Anscombe's argument, on one interpretation, is about our actual criteria for calling actions intentional.

We must distinguish the philosophic claim that actions are not called intentional in virtue of a certain feature from the socio-linguistic claim which might be expressed in just the same words. The philosophic claim is not a generalization over any actual events, either the occurrence of words or, say, the conscious applications of certain tests or criteria, whereas the socio-linguistic claim would be a generalization about what English-speakers do while or before calling an action intentional. For example, women are called "women" (in the sociolinguistic case) in virtue of certain features of dress, general shape, length of hair, and physiognomy, but they are called "women" (in the philosophic case) in virtue of anatomical features quite inaccessible in the normal case to the perception of either the language user or the observing sociolinguist. These ultimate criteria cannot easily be given a *locus existendi* since their actual use as criteria is reserved for bizarre cases, and sheltered individuals might use and understand the word "woman" perfectly well without even knowing that there were these ultimate criteria. It is true in general of philosophically puzzling ordinary words that although we all know how to use them and understand them, we must go to some effort to figure out the criteria for their use, and this effort does not involve a statistical survey, nor does the result merely remind us of what we have all along known. The philosophic claim in this case is not about what we often, sometimes, or always do when we call an action intentional, and hence the "everyday" mode of expression of the claim obscures the fact that it is essentially a schematic point, generated to be sure from the empirical facts of actual usage, but not a generalization over these facts.²

It might seem that the philosophic claim is normative, that "We do not call an action 'intentional', etc." is akin to the mother's admonition "We do not crayon on the walls." It is true that the philosophic claim is the expression of something like a rule, but it does not, as it were, enjoin us to obey it. The justification for the claim is not a giving of *reasons for* the existence of the rule (which are reasons for obeying it),

² This distinction has been variously expressed, and abused, by the ordinary language philosophers, often in rebuttal of Fodor and Katz's plea for a rigorous science of ordinary language philosophy based on field linguistics.

but a giving of *evidence* of its existence. Because the justification of the claim involves empirical evidence about usage, it is easily confused with the sociolinguistic claim, which is justified by the same *kind* of evidence. (The distinction in justification can be illuminated by the contrast between the evidence needed to confirm the existence of the rule in baseball: three strikes and you're out (supposing there is no rule book available), and the evidence needed to confirm the generalization: baseball players throw their bats whenever they have struck out.)

III

We can now examine Anscombe's conclusion that actions are not called intentional in virtue of a certain kind of feature. One interpretation of this is that it is a claim about our actual criteria, but not in the sense of "actual" that demands that we actually do apply them on every or indeed any occasion. In the light of this interpretation, the argument presented has the following form: Anscombe provides evidence for the existence of certain rather *roundabout* criteria for intentional actions — criteria dealing with further consequences, contexts, question and answer language games, etc. — and then sets out to show that no more *direct* criterion (such as a feature present in the person) operates, using the roundabout criteria already established as the standard by which to measure new candidates. Now since the philosophic claim about criteria is not about criteria always or ever consciously applied, any feature that *could* be checked for in a conscious way to determine whether a particular intentional action has occurred would have to count as a feature in virtue of which actions are called intentional.

How does "could" function in this instance? Any feature of the world which is fairly readily observable or ordinarily determinable by the ordinary people who make up the linguistic community *could* be used as the feature in virtue of which a certain word was used. Now if observation of such a feature would enable the language users to distinguish particular intentional actions, for example, then regardless of whether anybody in the community ever consciously applied the test of this feature before reaching a conclusion about intentions, this feature would be a feature in virtue of which actions are called intentional. Here we must be careful not to confuse symptom with criterion. Such a feature would be a feature in virtue of which etc., only if determining its presence in the particular case could *settle* the matter of intention, and this not on the basis of any contingent generalization or theory held to be true by the testers relating the presence of the feature normally with the meeting of other ("actual") criteria. That is, the claim that there is no

such feature is the claim that there is no fairly simple (but as yet undiscovered) question we can ask and answer about a present feature of a man that would settle the matter of his intentions.

How does one discover that this is the case? Very simply, it would seem. One conjures up a scene in which there is some doubt about a man's intention — e.g., he says "And now I push button A," and then pushes button B — and hunts about for some feature that would give away what the man intended — did he make a verbal slip, or a manual one, or neither? What is then quite obvious is that no one could just say, out of the blue, "Ah, did you notice that blink (or that swing of the arm, or any other feature)? That settles it: his intention was to confuse us into thinking he pushed A, while he surreptitiously pushed B." But does this obvious fact support the claim that "nothing about the man . . . can possibly determine the content of that description?" What if one of the observers had a theory about brain states, and announced that his electroencephalograph readings clearly showed what the man's intentions were? This is the watershed separating the modest from the ambitious interpretation of Anscombe's argument and conclusion. The modest argument is about what our *actual* criteria for intentional actions are, and the appropriate response to the brain-theory suggestion is that although one might be able to detect some feature of an actor that could serve, in a theory of motivation detection for example, as a *symptom* of intentional action, such a feature could only be a symptom, and never the criterion for calling an action intentional. If one *made* the feature the criterion, one would have *changed concepts*, and so would not be talking about intentional action (what *we* mean by intentional action) at all. So the modest form of Anscombe's argument says nothing about the impossibility of there being even a superreliable symptom, and merely makes the almost trivial point that we have certain roundabout ways of determining people's intentions and cannot tell what a man's intentions are by just looking at him.

Perhaps this is all Anscombe intended to show. If so, we should read her conclusion: actions are not *called* intentional (now) in virtue of a certain sort of feature. And the supporting claim should be read: nothing (readily observable) about the man *does* determine the content of the crucial description. But in arguing that the feature *I* must have certain characteristics, and no feature can have these characteristics, she seems to be going after bigger game, for this argument should hold of any feature, regardless of whether it is considered to be symptom or criterion. And her final statement of the conclusion is that "the assumption that some feature of the moment of acting *constitutes* (my emphasis) actions

as intentional leads us into inextricable confusions, and we must give it up." This surely suggests the ambitious argument.

The ambitious argument, then, is that since nothing can *possibly* determine the content of the action description, we could not discover a superreliable (or strongly reliable) symptom of intentional action, and hence even with a weak change of concept, there is no feature in virtue of which we call actions "intentional" (i.e., weakly changed "intentional"). This is the argument, as we shall see, that is needed if the divorce of action from the world of physical causes is to be more than a bit of trivial wordmongering. But now is the supporting unargued claim that nothing can possibly determine the content of the action description even plausible?

The ambitious argument can be seen to be parallel in certain respects to a quaint argument about magnets: no physical feature could be such that it would enable the object of which it was a feature to single out iron from among all other elements as attractive. We can imagine Thales, who said the lodestone has a soul, making just this argument, and in fact it has a certain weird plausibility. How *could* any atomic feature of a bit of stuff be responsible for this amazing attractive power, especially considering the selectivity of the power? And yet we know that given enough theory about atoms and magnetic attraction and iron and so forth, it is quite possible and plausible for there to be such a feature. The similarity between this argument and Anscombe's ambitious argument is just this: both arguments claim on a priori or perhaps intuitive grounds that a (physical) feature with certain characteristics is impossible. In the case of the magnet argument our intuitions deceive us; could the same be true of the argument about intentions?

IV

Anscombe's specifications of the characteristics that feature *I* would have to have allow us to frame the success conditions for finding *I*. Suppose everyone's face was constantly breaking out in patterns of color blotches, and suppose further that some sharp-eyed and clever person claimed that he could determine whether or not an action was intentional on the basis of these blotches. The test of this claim would be that he could regularly, and without failure, ascribe intention to actions (uniquely described actions) without applying any of Anscombe's roundabout criteria. The roundabout criteria would, of course, be used to test his ascriptions.

Clearly, if he is to be able to distinguish "push button A" intentional actions from "push button B" intentional actions, and the case where

one intentionally wakes the baby while intentionally sawing the plank, from the case where waking the baby is unintentional (and the success conditions require this), he will have to have some system for classifying blotches, in fact something remarkably like a blotch language. Given no system, he would be unable to ascribe intentions he had not previously examined via the roundabout criteria. That is, his *general* success in ascribing intentions to actions must depend on his having a *generative* system of blotch interpretation. Lacking a specifiable system, he would be written off as telepathist, a lucky guesser or a fraud. Now Anscombe is saying that such an eventuality is strictly impossible, but why? It does not seem to me logically impossible that people should have worn their hearts on their sleeves in this fashion, nor that had they done this someone would be able to figure out the system. It does not even seem to me to be improbable that such *in principle* decipherable systems of physical changes should exist in the brains of human beings.³ Using such discovered features as criteria for intentional action would, then, involve changing concepts, but only in the weak sense — or perhaps in the strong sense, since such a system would probably allow decisions (because of its generative nature) in those cases where by Anscombe's criteria no decisions are possible: where, for example, there is just not enough evidence of the truthfulness of our agent-respondent, and the "wider history of the incident" gives us no clues.

Supposing such a system of events or states in the brain, would these states have any relevance to the occurrence of intentional actions? That is, even if features were discovered that would uniquely determine the content of action descriptions, could these features in any way *constitute* actions as intentional; could they have more than a coincidental role in the production or performance of intentional actions? Anscombe says, at one point:

Then it is a mere happy accident that an *I* relevant to the wider context and further consequences *ever* accompanies the preintentional movements in which a man performs a given intentional action.

Here she seems to be conceding that although a feature *I*, say a thought of the action under the description under which the action is intentional, might coincide with the occurrence of the preintentional motions, there could be no connection, aside from mere contiguity or simultaneity, between the thought and the preintentional motions, since nothing about

³ In fact, the systematic ambiguity and redundancy of neural signals, which are required by the brain if it is to achieve reliable transmission using unreliable components, coupled with the sheer quantity of components (approximately 10¹⁰ neurons), rule out the practical possibility of this deciphering.

the man, etc., can possibly determine the content of the crucial description. The principle lurking behind this claim seems to be that nothing, or no feature, could possibly both cause the preintentional motions and uniquely determine the action description under which the preintentional motions make up an intentional action. But this principle is at best dubious. Why could not brain states be both garden-variety causes of preintentional motions, and states that, under a systematic interpretation, yield unique determinations of action descriptions? Consider the case of the "commands" sent to robot space vehicles or to the computation sections of computers. Are these not events of a system such that they are both causes of certain mechanical motions and under a systematic interpretation yield something like action descriptions ("Find Canopus," "Multiply x by y ")?

But, it may be argued, if a feature is the *cause* of the preintentional motions it cannot be the intention, or the feature constituting an action as intentional, for intentions are not causes, and intentional actions are not caused. Thus, Miss Anscombe argues: "Does (the actor) then notice that *I* is followed often enough by its description's coming true, and so summon up *I*? But that turns the summoning up of *I* into an intentional action itself, for which we shall have to look for a second *I*." That is, were *I* the *cause* of the action, then if the action is intentional, one must intend for *I* to occur, and this generates an infinite regress of causes and intentional actions. But why should my action not be called intentional if *I* — a physical feature — is caused just like any other physical feature, and is in turn the cause of my preintentional actions?

The standard answer is that intentional actions are not caused; part of what we mean when we say that an action is intentional is that it is not caused. The support for this claim refers us to an "ordinary" question and answer sequence:

- (1) Were you caused to do that (e.g., spill the coffee)?
- (2) No, I did it intentionally.

I doubt that this is ordinary. I doubt that anyone would ever speak just this way. But even if it is granted that we do talk in this way, what of it? How do I *know* I was not caused to spill the coffee? Have I non-inferential, nonobservational knowledge that I was *not* caused to do it? Or would it be better to say that I have inferential knowledge that at least certain *sorts* of causes were absent? That is, I know I didn't feel anyone bump my arm; I know I am not an epileptic, and I know that I just had the malicious thought "Let's make a mess of Smith's carpet." The question (1) is asking if anything like a bump or a twitch or a startling sight caused me to spill the coffee, and as far as I know, nothing like that did cause me to do it. It would be absurd to suppose

that when one asked (1) he intended to cover all physical and metaphysical eventualities with regard to causes, and that (2) is anything like a firm assertion of the absence of causes. Thus (1) is, most plausibly, an ellipsis for

(3) Did you do that because some external object or internal malfunction moved your body?

and (2) is an ellipsis for

(4) No cause of *that* sort operated — to the best of my knowledge. I did it intentionally (and I really have no idea what sort of causes if any *that* might involve).

This argument (and variations on it) is the only one I know which has the conclusion that intentional actions are not caused, and it seems to me to be a very poor one. There is another argument to the effect that intentions cannot be causes. Intentions cannot be causes, so the argument runs, because part of what we mean by "intending X" is that, "in the absence of interfering factors, it is followed by doing X."⁴ This being so, "intending X is followed, *ceteris paribus*, by doing X" cannot be a causal law, but a tautology. The relationship between intending and doing X is not contingent, and Hume says it must be if intending is the cause of doing. But by parity of reasoning conception cannot be the cause of pregnancy, since part of what we mean by conception is that *ceteris paribus* it brings on pregnancy. Since "conception" is *defined* as the cause of pregnancy, our "law" is a tautology, but that does not mean conception is not the cause of pregnancy. The point is that we can give other characterizations of the event we call conception, so that the contingency of our discovery is restored. Thus the argument that intentions cannot be causes must be that no independent characterization of intentions is *possible*, and for this argument to go through noncircularly it must not rely on the claim that for any other characterization of intentions proposed, this characterization cannot be a characterization of intentions, because that is not what we *mean* by intentions.⁵ As was the case with Anscombe's argument, the "modest" claim about criteria is trivial and insufficient.

V

The conclusions to be drawn from the examination of the ambitious argument are that the argument depends on the unargued claim that a certain sort of feature is impossible, and that this claim is implausible

⁴ Charles Taylor, *The Explanation of Behavior*, Routledge & Kegan Paul, and Humanities Press, 1964, p. 33.

⁵ Cf. Taylor, *op. cit.*, pp. 12-14, 45-48.

in spite of our initial intuitions. Moreover, the conclusion of the argument entails a further consequence which is at least remarkable, and to many minds unacceptable, namely, radical macroscopic physical indeterminism.

Whether or not actions are in any sense identical with physical motions of the body (and it seems strange to me to suppose that this makes sense), the occurrence of an action at least depends on the occurrence of motion, restricted in certain ways. Some action descriptions, such as "earning the esteem of one's peers" (which one can certainly intend to do), place no easily describable restrictions on the motions that might be said to fulfill this description, but other action descriptions, such as "putting one's right index finger in one's right ear" quite strictly limit the motions that would fulfill the description. "Mimicking Smith's gesture" and perhaps "dancing the role of the Firebird" put almost complete restrictions on the physical motions.

Now if there is no feature in virtue of which actions are intentional, then there is no *physical* feature in virtue of which actions are intentional. But if the occurrence of a particular action of mine is to be explained by its having been my intention to perform it, then the occurrence of the motions required to perform this action is to be explained with reference to no physical feature of me. That is, the antecedent conditions for the occurrence of these physical motions include a condition which is not a physical condition. Returning again to magnets, this claim about intentions is analogous to the claim that metal A at time t might be behaviorally magnetic, and at time t' not behaviorally magnetic, and the *only* difference between A at t and A at t' would be the truth of the governing conditional at t and the falsity of the conditional t' . That is, there might be no details of, say, submicroscopic alignment to account for or govern the magnetic disposition that would allow us to differentiate A at t from A at t' .

(There is a merely verbal escape from the clutches of this indeterminism which is perhaps worth mentioning. One might hold that the motions of my body (including preintentional motions) might be determined, and yet my actions not be determined, since actions are not identical with motions. Thus my arm might raise and my finger move the trigger as part of a deterministic universe, but I would retain the free choice of having this be an *action* of mine (presumably by acquiescing or annexing my *wanting* to the situation) or merely something that "happened to me." This is reminiscent of the confused Stoic position that free will consists in having the option of (internally) acquiescing in or fighting against one's determined actions.

It is important to note that the indeterminism entailed by this view

of intentional action is an indeterminism with regard to relatively macroscopic events. This is to be contrasted with a similar situation that occurs at the frontiers of particle physics, for example. For explanations of macroscopic events which cite some governing law or disposition, we are not prepared to entertain the possibility that we have reached an end to explanation, that the truth of the governing law is all that there is about the case. Where Ryle speaks of "inference tickets," Quine speaks of "promissory notes."

"From the start of a scientific attitude anyway, (the dispositional term "soluble") has been a sort of promissory note which one might hope eventually to redeem in terms of an explicit account of the working mechanism. What kind of account of a mechanism might pass as explanatory depends somewhat, of course, upon the general situation in science at the time. Nowadays, in general, one settles for an account in terms of the arrangement and movement of molecules, or of small particles when finer texture is relevant."⁶

The promise of deeper analysis holds until we reach the edge of particle physics. Here particles are given "characteristics" or "features" (such as a certain charge) solely on the basis of the truth of a certain dispositional law for which there is a test, and there is no compulsion to entertain the possibility of a deeper analysis of these "features." There is no compulsion because there is no nomological irregularity in there being no deeper analysis; the present depth of analysis suffices for our conceptual scheme. A macroscopic end of explanation is, on the other hand, a nomological asymmetry very much to be avoided. Thus the ambitious claim that actions are not called intentional in virtue of a certain feature presents us with no logical impossibility (for it is not logically impossible for there to be *only* the truth of a dead-end law linking regularities); it presents us with a nomological asymmetry that would require a vast adjustment in our scheme of the physical universe.

VI

My conclusion, then, is that Miss Anscombe's argument, that an action is not called intentional in virtue of any extra feature which exists when it is performed, is either trivial and hence fails to support the new doctrine about action, or fails because it depends on the claim that a certain sort of feature is impossible, and this claim is unsupported and implausible. Furthermore, the conclusion entails a nomological asymmetry which, though not logically absurd, should be viewed as utterly unlikely.

⁶ W. V. O. Quine, "Necessary Truth," in *Ways of Paradox*, Random House, 1966, p. 52.

The value of her argument is that it does give the conditions for the explanation of intentional action in terms of "features," and although she is wrong to claim that no features could possibly meet these conditions, she does point out the direction we must take to provide this explanation. It is interesting that these conditions she adduces are so in harmony with, and so suggestive of, the "information processing" models of the human brain. Carrying through with an "information processing" explanation of human action will not involve us in Identity Theory, as might be supposed. Since the systems of informational states, and their physical realizations, are bound to differ from brain to brain, no particular intention would be identified with a particular physical state, but rather it would be identified with similar *functional* states of different systems having different realizations.⁷ Thus causation of intentional actions (or if you wish, of the preintentional motions that fulfill but are not identical with intentional actions) is reintroduced without committing us to Identity Theory, or the view that reasons *are* causes.

Nomological symmetry is preserved, but the old problem of reconciling determinism with responsibility, which seemed to be so easily spirited away by the new doctrine about action, returns in slightly new guise and demands a more satisfactory answer.

D. C. DENNETT.

UNIVERSITY OF CALIFORNIA (AT IRVINE).

⁷ Hilary Putnam develops this concept of functional states in an unpublished paper, "The Nature of Mental States."