J. DAVID VELLEMAN

PRÉCIS OF THE POSSIBILITY OF PRACTICAL REASON*

Imagine that your arm becomes paralyzed but that your doctors predict a return to normal within a matter of days. When you wake up each morning, the first thing you do is to check whether you have regained control of your arm. What exactly are you hoping to find?

Part of what you are hoping to find, no doubt, is that your arm moves. But movement by itself would not be enough. Waking up to find your arm flapping around aimlessly would not lead you to think that your control over it had been restored. You would have to conclude instead that paralysis had given way to a spasm.

What you are hoping to find, then, is that your arm not only moves but moves when and where you want it to. Yet movement in response to your desires would not be enough, either. You might of course be encouraged if you found your hand scratching an itch behind your ear; but if you subsequently found it grabbing food off someone else's plate, you would not necessarily be reassured by the reflection that you had indeed wanted what he was eating.

The problem in this case, we might be inclined to say, is that although you wanted the food on someone else's plate, you also wanted to follow the rules of etiquette, and so grabbing the food was not something that you wanted to do on balance or overall. Your having control of your arm would require that the arm do, not just something that you wanted, but rather what you wanted on balance.

Yet how do you tell what you want on balance? If you have ever cast a speculative glance at the uneaten french fries on someone else's plate, you will know that the contest between

Philosophical Studies **121:** 225–238, 2004.
© 2004 *Kluwer Academic Publishers. Printed in the Netherlands.*

appetite and etiquette can be close. Surely, appetite might win out before you had realized that it was the stronger; indeed, you might realize that it was the stronger only by seeing that it had won out, as evidenced by the movement of your arm. And then the thought that this movement reflected the balance of your motives would not convince you of your having regained control.

In each of these cases, I have indicated your lack of control by casting you in the role of a spectator, who "sees" or "finds" that his arm is moving. Many philosophers share the intuition that being a spectator is the diametrical opposite of being an agent. As Brian O'Shaughnessy puts it, "Common to all experiences of loss of agency is the sense of becoming a spectator of one's own actions".[1] In the words of another philosopher, "[I]t seems as though someone has moved into your body and pushed you off the playing field up into the grandstand to be a mere spectator of yourself".[2] This trope is almost universal in the philosophical literature about action and the will.[3]

That this contrast seems so natural ought, on reflection, to seem odd. Why is *doing* so often contrasted with *seeing*? Why is the opposite of "participant" so often "observer" rather than "abstainer" or "absentee"? If our relation to things that are not our doing is the relation of onlooker, what does that say about our relation to the things that we do?

This question is the starting point for one venerable line of thought in the philosophy of action: if the essence of passivity with respect to an event is witnessing it, then perhaps activity with respect to an event, which can hardly consist in blindness to it, consists rather in knowing about it in some other way. This line of thought inspired the work that many regard as having inaugurated the contemporary study of action, Elizabeth Anscombe's *Intention*.[4]

Anscombe argued that what is our doing, in the fullest sense, can be distinguished from mere happenings by virtue of being the object of a special kind of knowledge.[5] Anscombe called it "knowledge without observation", and although I think that her attempts to define observation and its absence were not

successful, I also think that what she had in mind is clear enough.

Anscombe pointed out that expressions of intention, such as "I am going to take a walk", are similar to predictions insofar as they are "indicative (descriptive, informatory)" (§2, p. 3). When all goes well, they express the speaker's knowledge of what he is doing or will do, which is "known by being the content of [his] intention" (§30, p. 53). Anscombe thus conceived of intention as a potentially knowledge-bearing state, expressed in potentially knowledge-conveying utterances. What distinguishes the knowledge embodied in intention from other sorts of knowledge, she said, is that it is practical knowledge, in the sense that it causes – rather than being caused by, or causally concomitant to – the facts that make it true (§48, p. 87). By "knowledge without observation", then, Anscombe simply meant knowledge that is productive rather than receptive of what is known.

I have already suggested one reason for thinking that agency has something to do with a special kind of knowledge – namely, the familiar contrast between doing and seeing. Anscombe alluded to another reason by attributing the phrase "practical knowledge" to Aquinas, for whom it described God's knowledge of His creation. God knows what the world is like, but not by dint of having found out; He knows what the world is like because it is just as He meant it to be. And His meaning it to be that way already constituted knowledge on His part of how it would be, or rather how it already was. What's more, this epistemological relation that God bears to the world – knowing how it is just by meaning it to be that way – is constitutive of his role as the world's inventor or designer. The designer of something is the one whose conception of the thing determines how it is, rather than *vice versa*, and determines this not by chance but by a mechanism reliable enough to justify his confidence in that conception as an accurate representation of the thing. To be the designer of something is just to be the one whose conception of it has epistemic authority by virtue of being its cause rather than its concomitant or effect.

Anscombe's nod to medieval theology as her source for the term "practical knowledge" suggests that she conceived of intentional action as a realm in which human beings exercise a minor share of divinity. We invent our intentional actions, just as God invented the world, and our inventing them consists in our framing a conception of them that has epistemic authority by virtue of being determinative of them. Hence intentional action is behavior that realizes the agent's knowledge of it, just as the creation realizes God's omniscience.[6]

What I presume to call my theory of agency is in fact a variation on this theme of Anscombe's. In this section I will explain where my theory varies from Anscombe's, but I should first warn the reader that this theory is not the main topic of the book under discussion in the present symposium. I introduce it here because it is often the target of the discussion that follows. After describing the theory, I will turn to a description of the book.

I depart from Anscombe in two respects that are somewhat incidental to the philosophy of action. To begin with, Anscombe was – or, at least, tried to be – a thoroughgoing behaviorist about mental states. For example, in describing intentional action as that which gives application to a particular sense of the question "Why?", she tried to identify an overt language-game in terms of which intention could be understood. The problem is that she was also committed to the notion of practical knowledge, which requires intention to play the role of cause. Her obscure discussion of "mental causation" is a symptom of the resulting difficulties. I depart from Anscombe, then, in rejecting her behaviorism.

Anscombe was also, as I interpret her, a reliabilist about knowledge – in particular, about what is "known by being the content of intention". She thought that a reliable connection in general between what's intended and what's done is sufficient to confer the status of knowledge on a particular intention, provided that the connection holds up in the particular case. When I wrote my first book in the philosophy of action, I did not understand this aspect of Anscombe's view, and I consequently complained that she failed to provide any justification for the

knowledge that she described as "without observation". I still believe, as I argued there, that knowledge without observation can meet the justificatory standards of internalist epistemology; but I no longer think that this difference of opinion with Anscombe is especially significant.

My main departure from Anscombe has been to introduce a story about the dynamics of practical knowledge, based on two premises that seem to me uncontroversial. The first premise is that knowledge is not passively received in the manner imagined by some empiricists; it is the result of intellectual activity directed at a goal, which is the attainment of knowledge and understanding. Intellectual activity directed at that goal must have some motive force behind it, and that force must be guided by the intellect, not only with respect to means towards the goal, but also with respect to its very constitution – that is, with respect to what constitutes knowledge and understanding. Hence reason cannot just be the slave of the passions, as Hume believed; it must have some motives in its service, which might be called intellectual passions.

My second premise is that a normal person is aware, from his own egocentric perspective, of being identical with an especially salient member of the objective order – identical, that is, with the creature walking in his shoes, sleeping in his bed, eating his meals, and so on.[7] That creature is certainly of great interest to a person, and its doings consequently become the object of the person's intellectual motives. But the person's awareness of being identical with that creature opens up an obvious shortcut to the cognitive goal. The subject can know what that creature is doing simply by doing what he conceives of the creature as doing, or being about to do, since his conception will then turn out to be not only true but also justified, on the grounds of the creature's having this very incentive to bear it out. A person's conception of what he is doing has epistemic authority because he tends to behave in accordance with it; and he tends to behave in accordance with it so as to have, embodied in it, an epistemically authoritative conception of what he is doing. Practical knowledge thus supplants theoretical knowledge, as a more secure route to the same cognitive goal.[8]

There are many objections to the way that I use this story to extend Anscombe's conception of intentional action. Some of those objections are raised in this symposium and will be addressed in my answers to critics. But I am unaware of any cogent objection to the story itself or to the claim that it describes an irresistible route to self-knowledge. As it turns out, the psychological basis of the story has been copiously documented by social psychologists working in the area that is sometimes labeled self-consistency, an area that encompasses the topics of cognitive dissonance and attribution.[9] Research in this area, widely replicated over the course of decades, has shown that people have a broad tendency to behave in ways that are consistent with their own conceptions of themselves – of how they behave in general and of what their motives are on a particular occasion. People can be made to behave angrily by being convinced that they are angry, the more angrily, the more angry they are convinced of being. Children are more likely to be tidy if told that they are tidy than if told that they should be. Extremely shy people do not act shyly if led to believe that the symptoms of their social anxiety are attributable to something other than shyness. And so on.

One team of researchers observed that their subjects' behavior can be influenced by the act-descriptions that they are antecedently primed to frame, as if they have a tendency to fulfill antecedently framed descriptions of their forthcoming actions. These researchers hypothesized that this tendency is the means by which people know what they are doing, and that it constitutes the mechanism of acting on an intention.[10] Social psychology has thus arrived independently at a dynamic version of Anscombe's thesis that intention embodies practical knowledge.

The same researchers claim to have shown, furthermore, that we ordinarily seek to identify our behavior at a "high" or "comprehensive" level, representing our underlying motives and ultimate goals. They describe this tendency as a "search for meaning in action"[11] or "a human inclination to be informed of what we are doing in the most integrative and general way available".[12] Here the empirical findings harmo-

nize with my dynamic version of Anscombe's theory in a further respect.

With a now famous example, Anscombe pointed out that an agent often knows what he is doing under a series of descriptions each of which incorporates the answer to the question "Why?" directed at the same action under the previous description in the series. Why is he moving his arm? Because he is pumping water. Why is he pumping water? Because he is replenishing the water supply. Why is he replenishing the water supply? Because he is poisoning the inhabitants of the building. Why is he poisoning the inhabitants? Because he is assassinating enemy agents. With the exception of the first, purely physical description, all of the descriptions under which this person knows what he's doing are answers to the question why he is doing it as previously described.

The sequence from "moving his arm" to "killing enemy agents" displays a progression toward increasingly "high-level" or "comprehensive" act-descriptions. Hence if there is empirical evidence of "a human inclination to be informed of what we are doing in the most integrative and general way available", then it is evidence of an inclination to progress from rudimentary descriptions like the former towards comprehensive descriptions like the latter.

I believe that the existence of such an inclination follows from the two premises with which I have supplemented Anscombe's theory. The goal toward which our cognitive processes are directed must be, not merely registering rudimentary, observable facts, but also formulating them in "integrative and general" terms of the sort that convey understanding. When directed at our own behavior, these processes must be oriented toward the goal of knowing what we are doing in the sort of comprehensive terms that indicate why we are doing it, by alluding to the relevant dispositions and circumstances. And the previously described shortcut to self-knowledge – the shortcut of doing what we think we are doing, or are about to do – is also a route to this "high level" self-knowledge. We can attain integrative knowledge of what we are doing simply by framing and fulfilling integrative conceptions of our own

behavior, conceptions formulated in terms of the dispositions and circumstances that help to explain it.

In order to frame and fulfill integrative conceptions of our behavior, of course, we must be aware of a context with which to integrate it – projects and motives that we have, emotions that we feel, customs and policies that we follow, traits of character that we display, all of which afford terms for understanding our behavior as more than mere bodily movement. These other aspects of our self-conception – projects, motives, emotions, customs, policies, traits of character – will provide the materials for integrative knowledge of what we are doing, provided that we do things appropriately integrated with them. The goal of a more comprehensive knowledge of what we are doing therefore militates in favor of doing things that can be understood as motivated by our desires, expressive of our emotions, implementing our policies, manifesting our characters, and so on.

Here is the point at which I diverge from Anscombe most sharply. Those aspects of ourselves and our circumstances which we could incorporate into an integrative conception of doing something turn out to coincide with what we ordinarily count as reasons for doing it. When giving our reason for doing something, we often cite a desire that motivated it, an intention or policy that guided it, an emotion or opinion that animated it, a habit or trait that was manifested in it, and so on. Examples of desire-based reasons are well known; here are some examples of reasons based on other considerations that provide an explanatory context for an action:

Why are you whistling?
Because I'm happy.

Why aren't you having any wine?
Because I don't drink.

Why worry about his problems?
Because I'm his friend.

Why are you shaking your head?
Because I think you're wrong.

Why do you have her picture on your wall?

Because I admire her.

Here already?
I'm punctual.

Accordingly, I believe that reasons for doing something are facts that would provide an integrative knowledge of what we were doing, if we did that thing. Our cognitive processes will then favor framing and fulfilling a conception of ourselves as doing that thing, understood in the context of those facts, rather than other things for which we lack the elements of an equally integrative conception. As I have sometimes put it, reasons for doing something provide a *rationale*, an account in which our doing it is seen to cohere with our psyches and our circumstances. Whereas acting on an intention is a matter of realizing practical knowledge of what we are doing, acting for a reason is a matter of realizing more integrative practical knowledge, incorporating relevant facts that constitute reasons for acting.

My first book, *Practical Reflection*, explored the dynamics of practical knowledge in detail. Much of the book was devoted to drawing out the psychological and epistemological consequences of the hypothesis that framing and fulfilling a conception of what we are doing is an irresistible route to the goals of intellectual activity directed at our own behavior. Having drawn out these consequences, the book argued that they help to explain various aspects of human agency.

In the papers that I wrote immediately after the book – the papers collected in the volume under discussion in this symposium – I tried to step back from the details of the theory and focus on its overall shape. This goal is expressed in the Preface to the collection as follows: "Without even referring to the theory, I have tried to unearth more fundamental reasons for wanting a theory of its general form – reasons for thinking that there ought to be a theory of its kind" (vii).

The kind of theory for which I argue in two of the papers is one that adopts Anscombe's conception of intention as a state that potentially embodies knowledge. Leaving aside the details of my dynamic hypothesis about how such knowledge is attained, I try to show that an epistemic conception of intention

enables us to explain how the future can be open from the deliberative perspective ("Epistemic Freedom") and how intentions can be shared between two agents ("How to Share an Intention").

In another paper ("The Story of Rational Action"), I argue for my conception of reasons for acting as considerations that would give the agent an integrative knowledge of what he was doing, if he did that for which they are reasons. The basis of my argument in this paper is the normative force generally attributed to the axioms of formal decision theory. The fundamental theorem derived from those axioms states that if an agent's preferences satisfy them, then there will be utility and probability functions, unique within a linear transformation, according to which those preferences maximize expected utility. I argue that the theorem itself accounts for the normative force of the axioms, because it shows that preferences obedient to them will have a context into which they can be integrated – utility and probability functions in terms of which they can be comprehensively grasped and hence supported as by reasons for acting.

The three papers described thus far defend the kind of theory that results from taking seriously the notion of practical knowledge, as applied by Anscombe to intention, and as extended by myself to reasons for acting. In five other papers I defend the kind of theory that identifies a constitutive aim of action, an aim with respect to which behavior must be somehow regulated in order to qualify as action, in the same way that cognition must be regulated for truth in order to qualify as belief. The aim that I have in mind is the aim of intellectual activity directed at our own behavior – that is, the aim of having an integrative knowledge of what we are doing. But in general I try to abstract from this particular aim, in order to argue that there must be some aim or other by virtue of which behavior qualifies as action.

The first of these papers, "The Guise of the Good", argues that Davidson's theory of action fails to account for the normative force of reasons for acting. In Davidson's theory, the normative force of reasons depends on value judgments that are somehow implicit in desires or other conative attitudes. I argue

that even if desiring something entails regarding the thing as good in some sense, this evaluative regard lacks the affirmative force of a value judgment and is therefore unfit to serve as a premise in a practical syllogism, as Davidson thought it could. This paper motivates the search for an alternative explanation of how reasons for acting get their normative force, a search that leads, in two subsequent papers, to the idea of action's having a constitutive aim.

In "The Possibility of Practical Reason", I try to show that identifying a constitutive aim for action would break the impasse between so-called internalists and externalists about reasons, by tracing the force of reasons back to the agent's motivational makeup without making their content depend on the contingent inclinations in which he may differ from other agents. What lends reasons their force, according to this explanation, is an inclination without which a person is not an agent at all – is not in the business of acting – just as he is not in the business of believing unless his cognitions are regulated by an inclination toward the truth. In "Deciding How to Decide", I argue that this strategy of explanation is preferable to the strategy pursued by David Gauthier, of finding pragmatic justification for the norms of practical reasoning.

"What Happens When Someone Acts?" argues that conceiving of action as behavior regulated by some constitutive aim would yield an account of agent causation more satisfactory than either the causal theory of Davidson or the hierarchical theory of Frankfurt. Finally, the paper "On the Aim of Belief" defends the conception of belief after which I have modeled my conception of action as having a constitutive aim.

In none of these papers do I say that behavior qualifies as action when it is regulated by the aim of knowing what we are doing. In the originally published version of one paper, "The Possibility of Practical Reason", I identified a different aim, autonomy, as constitutive of action; and I decided not to alter that paper when including it in the collection, despite having concluded that it was mistaken in that respect. Otherwise, however, the papers focus on reasons for wanting to identify a constitutive aim for action rather than on actually identifying one.

When gathering the papers into a volume, however, I decided to write an introductory overview, whose purpose is stated in the Preface as follows:

> The Introduction is an attempt to fashion a single narrative out of the main themes that appear in the rest of the collection. In concentrating on the flow of this narrative, I have tended to gloss over argumentative details, relying on the other papers to provide them. I have tried to indicate in the footnotes where detailed versions of the argument can be found in the other chapters. The Introduction also records recent changes of mind about various issues. [vii]

The primary change of mind that I wanted to record in the Introduction concerned the aim identified as constitutive of action in the title essay. Here is how I reported that change of mind:

> In the title essay, "The Possibility of Practical Reason", I identified the constitutive aim of action as autonomy itself, partly because I liked the Kantian ring of that claim and partly because I had hopes of forestalling criticism of my view as oddly intellectualist, or as portraying an autonomous agent to be unduly self-absorbed. Unfortunately, I think that the resulting version of my view is unworkable, as becomes evident, I think, in the final, tortured sections of that paper.

In order to correct the error noted here, I fashioned the introductory narrative to include my considered view about the constitutive aim of action – the view that it is the aim of our intellects as focused on ourselves, the aim to which practical knowledge is the obvious shortcut, the aim of knowing what we are doing. But that view is not directly defended in any of the papers collected in the volume.

In this respect, the book turned out to be flawed when considered as a self-contained monograph: the Introduction makes a claim that is not defended elsewhere in the volume. Of the commentators in this symposium, two focus on that flaw, while the third charitably reads the book as a fragment of a larger research program.

NOTES

* Thanks to Nishi Shah for extensive comments on my contributions to this symposium.

[1] *The Will*; *A Dual Aspect Theory*. Cambridge: Cambridge University Press, 1980, Vol. 2, p. 36.
[2] Kubara (1975).
[3] Al Mele (1997, p. 238) uses this image in his "Agency and Mental Action". Other uses include: "[O]ne's capacity to govern one's conduct is undermined … . One is reduced to a spectator" (Audi, 1986, p. 534); "The sense of autonomy is the sense that one is not merely a witness to one's life but rather fashions it from the world as one finds it (Watson, 2003); "The agent must not be a mere bystander or onlooker of what happens" (Pettit, 2001); "This makes it seem as if getting up from a chair is something that happens to a man, something to which he is at best a spectator" (Dilman, 1999); "I am not just a spectator of my life, but the real actor in it" (Honderich, 2002); "One is not normally in a *passive* relationship with such features of one's behavior, and is an agent who deliberates, decides, and acts out one's decisions, not a spectator of forces carrying one along" (Smilansky, 2002); "Thinking is not something that occurs to you, like the beating of your heart, something concerning which you are a mere spectator. Thinking is something you do" (Heil, 1998); "[I]n the experience of deliberation, we are not mere spectators of a scene in which … contending desires struggle for mastery with ourselves as the prize" (Strawson, 1992); "When you determine yourself to be the cause of your action you must identify yourself with the principle of choice on which you act. … In this kind of case, you do not regard yourself as a mere passive spectator …" (Christine Korsgaard, "Practical Reasoning and the Unity of the Will", lecture II in the John Locke Lectures).
[4] 2nd edn. Cambridge, MA: Harvard University Press, 2000.
[5] I discuss Ancombe's *Intention* more fully in "What Good is a Will?" (MS). I also discussed it in my *Practical Reflection* (Princeton: 1989; http://www.umich.edu/~velleman/Practical_Reflection/), but I now regard that earlier discussion as based on an inadequate understanding of Anscombe's views.
[6] Here "intentional action" is Anscombe's term. My own view is that the category of intentional action is not a natural kind of behavior, because its boundaries are determined in part by norms of moral responsibility. A discussion of God's knowledge as practical in the same sense appears in Maimonides' *Guide for the perplexed*, part III chapter 21.
[7] I discuss this point at length in "The Centered Self" (MS).
[8] This point was a theme of Hampshire (1959).
[9] For a discussion of this research, see my "From Self-Psychology to Moral Philosophy", 2000. Some of this material is summarized in footnotes 32 and 33 on p. 27 of my book. Note that I do not regard empirical evidence for this story as essential to its philosophical interest. To my mind, it is a story of how agency might be realized, whether or not it is the story of how agency is realized in fact.

[10] See the publications of Wegner and Vallacher cited in the following two notes.
[11] Wegner and Vallacher (1986).
[12] Vallacher and Wegner (1985).

REFERENCES

Al Mele (1997): 'Agency and Mental Action', *Mind, Causation and World, Philosophical Perspectives* 11, 231–249.

Audi, R. (1986): 'Acting for a Reason', *Philosophical Review* 95, 511–546.

Dilman, I. (1999): *Free Will; An Historical and Philosophical Introduction*, London: Routledge, p. 264.

Hampshire, S. (1959): *Thought and Action*, London: Chatto and Windus.

Heil, J. (1998): *Philosophy of Mind; A Contemporary Introduction*, London: Routledge, p. 73.

Honderich, T. (2002): *How Free Are You? The Determinism Problem*, Oxford: Oxford University Press.

Kubara, M. (1975): 'Acrasia, Human Agency and Normative Psychology', *Canadian Journal of Philosophy* 5, 215–232.

Pettit, P. (2001): *A Theory of Freedom; From the Psychology to the Politics of Agency*, New York: Oxford University Press, p. 10.

Smilansky, S. (2002): *Free Will and Illusion*, Oxford: Oxford University Press, p. 207.

Strawson, P.F. (1992): *Analysis and Metaphysics: An Introduction to Philosophy*, Oxford: Oxford University Press, pp. 134–135.

Vallacher and Wegner (1985): *The Theory of Action Identification*, Hillsdale, NJ: Erlbaum, p. 26.

Velleman, J.D. (2000): 'From Self-Psychology to Moral Philosophy', *Philosophical Perspective* 14: *Action Theory and Freedom*, 349–377.

Watson, G. (2003): "*Introduction*" *to Free Will*, Oxford: Oxford University Press, p. 1.

Wegner and Vallacher (1986): 'Action Identification', in R.M. Sorrentino and E.T. Higgins (eds.), *Handbook of Motivation and Cognition*, New York: Guilford Press, pp. 555–566.

Department of Philosophy
University of Michigan
435 South State Street
Ann Arbor, MI 48109-1003
USA
E-mail: velleman@umich.edu