



Philosophical Review

Two Faces of Intention

Author(s): Michael Bratman

Source: *The Philosophical Review*, Vol. 93, No. 3 (Jul., 1984), pp. 375-405

Published by: [Duke University Press](#) on behalf of [Philosophical Review](#)

Stable URL: <http://www.jstor.org/stable/2184542>

Accessed: 16/11/2010 11:59

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=duke>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Duke University Press and *Philosophical Review* are collaborating with JSTOR to digitize, preserve and extend access to *The Philosophical Review*.

<http://www.jstor.org>

TWO FACES OF INTENTION

Michael Bratman

We do things *intentionally*, and we *intend* to do things. Our commonsense psychology uses the notion of intention to characterize both our *actions* and our *mental states*: I might intentionally start my car, and I might intend to start it. My intention to start it clearly does not guarantee that I do. But beyond that it is not obvious how these two phenomena *are* related. A central problem for a theory of intention is to provide a plausible account of this relation.

One thing seems clear: it is part of our commonsense psychological framework that these phenomena are not completely *unrelated*. In classifying both our actions and our states of mind in terms of some root notion of intention, commonsense psychology clearly assumes that there is some important commonality. Our problem is to say what this commonality is, by spelling out the relation between intentional action and intending (or, having an intention) to act.

There are two common approaches to this problem. The first—the *desire-belief model*—sees intentional action as action that stands in appropriate relations to the agent's desires and beliefs.¹ This is a reductive model: it sees intentions to act as reducible to certain desires and beliefs.² On this approach, the problem of the relation

¹This is one of the guiding themes in Davidson's classic paper, "Actions, Reasons and Causes," in *Essays on Actions and Events* (New York: Oxford University Press, 1980). See also A. Goldman, *A Theory of Human Action* (Englewood Cliffs: Prentice-Hall, 1970). Davidson's commitment to this tradition has waned in recent years. I discuss this change in Davidson's theory of action, and its relation to some of the issues discussed in the present paper, in my "Davidson's Theory of Intention," forthcoming in a Festschrift in honor of Donald Davidson, edited by Merrill Hintikka and Bruce Vermazen, (Oxford: Oxford University Press, 1984).

²Philosophers who explicitly defend a reduction of the state of intention to certain desires and beliefs include R. Audi, "Intending," *Journal of Philosophy* 70 (1973), pp. 387–403; M. Beardsley, "Intending," in A. Goldman and J. Kim, eds., *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson and Richard Brandt* (Dordrecht: Reidel, 1978); and P. Churchland, "The Logical Character of Action-Explanations," *The Philosophical Review* 79 (1970), pp. 214–236.

between acting intentionally and having an intention to act becomes the problem of the relation of the complex of desires and beliefs constitutive of the latter to those desires and beliefs necessary for the former.

I think this approach is mistaken. We are planning creatures. We frequently settle in advance on plans for the future. On occasion, this even involves settling on one of several conflicting options each of which is, in light of our desires and beliefs, equally attractive. These plans help guide our later conduct and coordinate our activities over time, in ways in which our ordinary desires and beliefs do not. Intentions are typically elements in such coordinating plans. Once we recognize this central role intentions play in our lives the natural view to take, I think, is that intentions are distinctive states of mind, not to be reduced to clusters of desires and beliefs.

So I have argued in several recent papers.³ Here my argument against the desire-belief model will be only indirect. I will try to show what a part of a theory of intention would look like once we reject that model. Insofar as the account sketched is plausible, it will constitute an indirect argument against that model.

This brings us back to our problem of the relation between intending to act and acting intentionally. Once we see intentions as distinctive phenomena, how should we understand this relation? Here is where the second common approach comes in. I may intend to start my car later today: this is a *future-directed* intention. But I may also intend to start my car now: this is a *present-directed* intention. Such a present-directed intention does not guarantee that I actually start my car. But if I do start my car intentionally then, it seems plausible to suppose, I have such a present-directed intention to start it. After all, while starting the car I surely intend to do *something*. Given that what I do intentionally is start it, it seems that what I intend will include starting it.

³"Intention and Means-end Reasoning," *The Philosophical Review* 90 (1981), pp. 252–265; "Taking Plans Seriously," *Social Theory and Practice* 9 (1983), pp. 271–287; "Davidson's Theory of Intention," op. cit. I discuss slightly different, but related matters in "Castañeda's Theory of Thought and Action," in James Tomberlin, ed., *Agent, Language and the Structure of the World: Essays Presented to Hector-Neri Castañeda, with His Replies* (Indianapolis: Hackett, 1983).

This suggests a general solution to our problem: for me intentionally to *A* I must intend to *A*; my mental states at the time of action must be such that *A* is among those things I intend. I will call this the *Simple View*.⁴

The Simple View is a special case of a more general conception, the *Single Phenomenon View*. On this more general view, intentional action and the state of intention both involve a certain common state, and it is the relation of an action to this state that makes that action intentional. The Simple View adds to this more general

⁴Note that the Simple View does not say that there must be a separate *event* of intending to *A* for each intentional *A*-ing. The Simple View only imposes a requirement on one's mental *states*. Philosophers who accept something tantamount to the Simple View include B. Aune in *Reason and Action* (Dordrecht: Reidel, 1977), Chapter II, esp. pp. 89–102, and J. Searle in *Intentionality* (Cambridge: Cambridge University Press, 1983), Chapter 3. Searle says that the rejection of what I have called the Simple View is “a mistake that derives from a failure to see the difference between prior intentions [what I have called future-directed intentions] and intentions in action [what I have called present-directed intentions]” (p. 94n). But, as will be seen, my objection to the Simple View does not depend on such a failure.

In *The Principles of Morals and Legislation* (New York: Hafner, 1948), p. 84, Bentham famously distinguishes between consequences which are “directly” intentional and consequences which are only “obliquely” intentional. This distinction suggests a view according to which bringing about *X* may be intentional, even if one does not intend to bring *X* about, so long as one intends something one expects will (or, will likely) result in *X*. Such a view is intermediate between the Simple View and the more complex view I will be sketching below in Sections 4–6. This intermediate view is represented in contemporary discussions by, among others, H. N. Castañeda in his important book, *Thinking and Doing* (Dordrecht: Reidel, 1975), esp. p. 313. Here I will just note that this intermediate view is also subject to the objection against the Simple View to be developed in Sections 1 and 2.

The Simple View is rejected in passing by Georg Henrik von Wright in his *Explanation and Understanding* (Ithaca, N.Y.: Cornell University Press, 1971), pp. 89–90. More recently, Gilbert Harman has criticized the Simple View in “Practical Reasoning,” *Review of Metaphysics* 29 (1976), pp. 431–463. I discuss Harman's criticism, and my reasons for preferring my alternative approach, below. Despite this and other differences in our views, my thinking about intention has been influenced by this paper and by Harman's other two recent papers on this subject: “Willing and Intending,” forthcoming in *Philosophical Grounds of Rationality*, edited by R. Grandy and R. Warner (Oxford: Oxford University Press); and “Rational Action and the Extent of Intentions,” *Social Theory and Practice* 9 (1983), pp. 123–141.

conception the requirement that this state is just an intention so to act.

The Simple View has its virtues. It recognizes the distinctiveness of intentions, and provides a straightforward and initially plausible account of the relation between such intentions and intentional action. It is a view towards which commonsense initially leans, as well as a view implicit in many discussions of intention in moral philosophy.⁵ Nevertheless, while I will be accepting a version of the more general, Single Phenomenon View, I find the Simple View unacceptable. Our conception of the state of intention is that of a single state tied to two very different sorts of phenomena. Intention is Janus-faced, tied both to coordinating plans and intentional action. The Simple View does not allow sufficient theoretical room for both these faces of intention.

In Sections 1. and 2. of this paper I explain why. In Section 3. I show how one might naturally be led to the Simple View by an unacceptable reduction of another kind: a reduction of present-directed intention to volition. Finally, in Sections 4. through 6. I sketch a route between the desire-belief model and the Simple View, a route that remains within the framework of the more general Single Phenomenon View. My proposal sees intentions as distinctive, and sees the intentionality of an action as dependent on its relation to such intentions. But it rejects the account of this relation provided by the Simple View. It holds, instead, that while to *A* intentionally I must intend to do *something*, I need not intend to do *A*. This leads to a distinction between *what* I intend and the *motivational potential* of my intention. I conclude by arguing that this distinction has a further virtue: it allows our concern with the ascription of responsibility to shape our classification of actions as intentional without thereby distorting our classifications of mental states in ways which undermine critical regularities.

1. CONSISTENCY OF INTENTION AND THE SIMPLE VIEW

My argument against the Simple View is rooted in my conception of intentions as elements in coordinating plans. So I need to say more about that conception.

⁵See, for example, Charles Fried, *Right and Wrong* (Cambridge, Mass.: Harvard University Press, 1978), esp. pp. 20–24.

We have been speaking of present-directed intentions. But there is a tension in saying that I intend to do what I am *now* doing: talk of what I intend to do seems normally reserved for my attitude towards my *future* conduct. When I am actually starting my car it may seem natural to say that I no longer *intend* to start it, I *am* starting it. I think we should take this strain as a philosophical hint: not that there are no present-directed intentions, but that to understand what intentions are we should begin by concentrating on the future-directed case.⁶ This is the *methodological priority of future-directed intention*.⁷

Future-directed intentions are typically elements in larger plans.⁸ Such plans help me to coordinate my activities over time, and my activities with yours. The ability to settle in advance on such

⁶Note that even my present-directed intention to start my car is an intention to perform an action that continues somewhat into the future. Indeed, I doubt whether it is possible to have a present-directed intention to perform an instantaneous action, for reasons outlined by Brian O'Shaughnessy in *The Will* (Cambridge: Cambridge University Press, 1980), vol. 2, pp. 312–313.

⁷Here I diverge from a long-standing tradition in the philosophy of action. This tradition begins with Anscombe's decision in her groundbreaking monograph to treat intentional action, rather than intending to act, as the basic case in terms of which to understand intention. See *Intention* (Ithaca, N.Y.: Cornell University Press, 1963), esp. p. 9.

⁸There is a certain ambiguity in talk about plans. Sometimes we are talking about *states* of the agent—states of *having* certain plans. Other times we are talking about an appropriate abstract structure—some sort of partial function from circumstances to actions, perhaps—that may be used to describe the planning-states of different people. A more careful usage might reserve 'plan' for the latter and 'having a plan' for the former; but this is frequently stylistically awkward. In this essay I use 'plan' to mean 'having a plan'—that is, a state of mind. Thus plans are in the same category as (though different from) desires and beliefs.

It is worth noting that the importance of plans, and the associated phenomenon of planning, to our understanding of intention is sometimes blocked from view simply by the terminology we employ. For example, in his paper, "Intention and Punishment," in *Punishment and Responsibility* (Oxford: Oxford University Press, 1968), H. L. A. Hart writes

Intention is to be divided into three related parts. . . . The first I shall call 'intentionally doing something'; the second 'doing something with a further intention', and the third 'bare intention' because it is the case of intending to do something in the future without doing anything to execute this intention now (p. 117).

This scheme forces us to see most future-directed intentions merely as "bare intentions," and this tends to block from view the roles of such intentions in plans, and the resulting constraints on these intentions.

plans enables us to achieve complex goals we would not otherwise be able to achieve. This ability to settle on coordinating plans is a kind of universal means: it is of significant use in the pursuit of goals of very different sorts.

Intentions aid coordination as elements in larger plans. The concern with coordination exerts pressure towards unification of our various intentions. So if our intentions are to be well-suited to aid coordination, we should be able to put them together into a larger plan which can serve this coordinating role well. But to coordinate my activities over time a plan should be, other things equal, internally consistent. Roughly, it should be possible for my entire plan to be realized.⁹ Further, a good coordinating plan is a plan for the world I find myself in. So, assuming my beliefs are consistent, such a plan should be consistent with my beliefs, other things equal. Roughly, it should be possible for my entire plan to be realized while my beliefs are true.¹⁰

Let us say that my intentions are *weakly consistent* if they could all be put together into an overall plan that is internally consistent. My intentions are *strongly consistent, relative to my beliefs* if all my intentions could be put together into an overall plan that is consistent with those beliefs. To be well-suited to aid coordination, my intentions will need to be, other things equal, strongly consistent relative to my beliefs. Since it is largely to aid such coordination that we bother with future-directed intentions in the first place, we have a pragmatic rationale for a rational demand that future-directed intentions be strongly consistent, relative to the agent's beliefs. This is a demand that should be respected in our further practical reasoning and planning.

This demand for strong consistency distinguishes intentions from ordinary desires. I might, without irrationality, both desire to play basketball today and desire to finish this paper today, all the

⁹I assume a broad notion of realization such that a conditional intention to *A* if *p* is realized if not-*p*.

¹⁰I do not think that these brief explanations of the relevant notions of consistency are without their difficulties. There are deep problems here analogous to problems that arise when we try to say, for example, in what sense our beliefs about the morning star (that is, the evening star), or Cicero (that is, Tully) should be consistent. But I think this gloss on consistency will suffice for my purposes here.

time knowing I cannot do both. In contrast, intentions to play and to finish would, given my beliefs, convict me of a criticizable form of irrationality.

The demand for strong consistency provides the basis of my argument against the Simple View. But first I need to make one more point about that view. Suppose I intentionally start my car. On the Simple View it follows that

(1) I intend to start my car.

The point to note is that I can have the intention reported in (1) whether or not I actually do start my car. As we might say, the form of (1) is *not*

(2) *aRb*

where *b* is replaced by a singular term denoting an actual, particular action of starting my car.¹¹

This clarification in mind, let us turn to a series of three examples.¹² In the first case I am playing a video game in which I am to guide a “missile” into a certain target. I am quite skilled at such things, but it is a difficult game and I am doubtful of success. Still, I aim at the target and try to hit it. As it happens, I succeed in just the way I was trying. My success was not merely a matter of luck; it depended heavily on my considerable skills at such games. Further, hitting the target was what I wanted to do; I was not just aiming at the target as a way of ensuring that the “missile” would go several inches to the right.¹³

Do I hit the target intentionally? It seems that I do. I want to hit it and so am trying to hit it. My attempt is guided by my perception of

¹¹This explains why I did not include Donald Davidson among those who accept the Simple View, even though he comes close to endorsing the view that if I intentionally start my car then I must intend my particular act of starting it. See his essays, “How Is Weakness of the Will Possible?” and “Intending,” both in *Essays on Actions and Events*, op. cit.

¹²These examples take off from an example sketched by Robert Audi, op. cit., esp. p. 401.

¹³In this last sentence I am indebted to Harman’s discussion of such examples in “Willing and Intending,” op. cit.

the target. I hit the target in the way I was trying, and in a way that depends on my relevant skills. And it is my perception that I have hit it that terminates my attempt. So even though I am doubtful of success while I am trying, if I do succeed in hitting the target I hit it intentionally. On the Simple View, then, I must intend to hit the target. And this is, for all we have said, an acceptable result.¹⁴ Even though I am doubtful that I will hit the target, the intention to hit it need not violate the demand for strong consistency.

Suppose now that a second game is added, a game which also involves guiding a "missile" into a certain target. Since I am ambidextrous and can play one game with each hand, I decide to play both games simultaneously. As before, the games are difficult and I am doubtful of success at either of them. As it happens, I miss target 2 but I do succeed in hitting target 1 in the way I was trying and in a way that depended on my relevant skills. Here again, it seems to me, I hit target 1 intentionally. The mere fact that I was also trying unsuccessfully to hit target 2 does not prevent me from intentionally hitting target 1.

The Simple View must say, then, that I intend to hit target 1. And this seems plausible. But what about my intentions concerning target 2? I was trying equally hard, and with equal skill, as well as with equally weak confidence of success, to hit target 2. It seems clear from the symmetry of the case that if I intend to hit target 1 I also intend to hit target 2. Of course, in the example I do *not* hit target 2, whereas I do hit target 1. But, as we noted above, this difference does not prevent me from intending to hit target 2.

So the defender of the Simple View must suppose that in this case I intend to hit each target. This sets the stage for my argument against this view, an argument which requires one more addition to our example.

Let us now suppose that the two games are known to me to be so linked that it is impossible to hit both targets. If both targets are about to be hit simultaneously the machines just shut down. Both targets remain visible to me, so I can see which target I hit if I hit either target. And there is a reward for hitting either target. But I

¹⁴One might even here object to the Simple View if one thought that to intend to hit the target I must believe I will. Below I discuss this line of argument against the Simple View, and why I do not take it.

know that while I can hit each target, I cannot hit both targets. Still, I know it is difficult to hit either target, so I again decide to play both games simultaneously; I see the risk of shutting down the machines as outweighed by the increase in my chances of hitting a target. I proceed to try to hit target 1 and also to try to hit target 2. I give each game a try.

Suppose I do hit target 1 in just the way I was trying to hit it, and in a way which depends heavily on my considerable skills at such games. It seems, again, that I hit target 1 *intentionally*. So, on the Simple View, I must intend to hit target 1. Given the symmetry of the case I must also intend to hit target 2. But given my knowledge that I cannot hit both targets, these two intentions fail to be strongly consistent. Having them would involve me in a criticizable form of irrationality. But it seems clear I need be guilty of no such irrationality: the strategy of giving each game a try seems perfectly reasonable. If I am guilty of no such irrationality I do not have both of these intentions. Since my relevant intentions in favor of hitting target 1 are the same as those in favor of hitting target 2, I have neither intention. So the Simple View is false. If it were true I would be guilty of a form of criticizable irrationality; but I need be guilty of no such irrationality. The Simple View imposes too strong a link between intention and intentional action, a link that is insensitive to differences in the demands of practical reason.

This argument against the Simple View appeals to constraints on intention that do not apply in the same way to intentional action. In this respect it is similar to an alternative argument that has been sketched in the literature. It will be useful to discuss this argument briefly.

Suppose I intend now to go to the concert tonight. What must I believe about my future concert-going? Some philosophers¹⁵ accept the strong thesis that I must now believe I will go. Their reasons for this strong thesis tend to be of two sorts. There is, first, the need to explain the apparent oddness of remarks like: "I intend to go to the concert, but I may not go."¹⁶ Second, there is

¹⁵For example, H. P. Grice, "Intention and Uncertainty," *Proceedings of the British Academy* 57 (1971), pp. 263–279; and G. Harman, "Practical Reasoning," *op. cit.*

¹⁶Grice, *op. cit.*, pp. 264–266.

the idea that by seeing intention in this way we can best explain the role of intentions in various kinds of practical thinking.¹⁷ I will not examine such arguments for this strong belief condition here. It suffices for my purposes to note that once we are given this strong belief requirement on future-directed intention it will be natural to suppose that present-directed intentions are subject to a similar belief condition; and this leads directly to an argument against the Simple View.¹⁸

This argument has two premises. The first is just this strong belief requirement. The second is the observation that a person can do something intentionally even though, at the time of action, he is in doubt whether he is so acting. We have already seen an example of this: I might intentionally hit the target even while being doubtful of success. Donald Davidson offers another example.¹⁹ A person might try hard to make ten carbon copies on a typewriter, while being skeptical of success. Still, if this is what he wants to do, and if he does, in fact, make ten copies in the way he was trying and in a way that depends on his relevant skills, then it seems that he intentionally makes ten copies. Again, we have intentional action despite lack of belief.

So we have two premises: a strong belief requirement on intending to act, and the observation that one may *A* intentionally even while doubting that one is *A*-ing. These two premises entail that the Simple View is false. Given the strong belief requirement, when I act intentionally in a way in which I do not believe I am acting I will not intend so to act.

Like my initial argument, the present argument tries to cite a constraint on intention that does not similarly apply to intentional action. But whereas I cited the constraint that rational intentions are to be strongly consistent, given the agent's beliefs, the present argument cites a strong belief condition on intention. Now it seems to me that this strong belief condition is problematic in ways in which the demand for strong consistency is not. It seems plausible to suppose that sometimes intentions just do not satisfy such a

¹⁷For example, one of Harman's arguments in favor of this strong thesis is that it allows for a natural account of the role of intentions in means-end reasoning. *Op. cit.*, p. 435.

¹⁸As Harman explicitly notes. *Ibid.*, p. 433.

¹⁹In "Intending," *op. cit.*

strong belief condition. For example, I might intend now to stop at the bookstore on the way home while knowing of my tendency towards absent-mindedness—especially once I get on my bike and go into “automatic pilot.” If I were to reflect on the matter I would be skeptical about my stopping there, for I know I may well forget. It is not that I believe I will not stop; I just do not believe I will. Still, my plan is to stop.

Examples like this seem at least to show that the strong belief requirement is no more obvious than the Simple View itself. So a philosopher committed to the Simple View could plausibly resist the present argument by turning it on its end and seeing it as an objection to the strong belief requirement. One person’s *modus ponens* is another’s *modus tollens*.²⁰

In contrast, the demand for strong consistency of intentions is more difficult to avoid. First, instead of requiring an actual belief that I will *A* for me to intend to *A*, it demands only that (other things equal, and if my intentions are to be rational) I not have beliefs inconsistent with the belief that I will *A*. Second, this constraint is even compatible with the possibility of my intending to stop at the bookstore and believing I will not. It just requires us to say that, other things equal, I would then be guilty of a form of criticizable irrationality. Finally, it will be more difficult to turn the tables on my argument, rejecting the requirement of strong consistency in order to hold onto the Simple View. This is because this consistency constraint seems to be firmly grounded in a basic feature of intentions: their role in coordinating plans.

Nevertheless, objections to my argument remain. I turn now to consider some of these.

2. OBJECTIONS AND REPLIES

My argument depends on two claims about the final video games case in which the games are known to be linked and I succeed in hitting target 1:

²⁰Of course, if we reject the strong belief requirement, we will need alternative explanations of the linguistic data and the data about practical reasoning originally summoned in its support. I sketch an alternative explanation of the latter in my “Intention and Means-End Reasoning,” *op. cit.* I think an alternative explanation of the former can also be constructed, though I will not try here.

- (i) If in this case I had present-directed intentions which failed to be strongly consistent, I would be criticizably irrational.

and

- (ii) I hit target 1 intentionally.

Let us consider some ways in which a defender of the Simple View might try to challenge these claims.

Begin with (i). It might be urged that, for all that I have said, only *future*-directed intentions are subject to the strong consistency requirement. So I can intend to hit target 1 *now*, and similarly concerning target 2, without being criticizably irrational, contrary to (i).

This objection is inadequate for two reasons. First, the argument for the demand for strong consistency depended on the observation that intentions typically play a coordinating role. Now, while this is clearest in the case of future-directed intentions, this is also an important role of some present-directed intentions. Suppose my intentions concerning the video games are embedded in a larger plan for the day. I begin the day with what are then future-directed intentions concerning these games. When the time comes these become present-directed intentions. But they continue to be part of my coordinating plan. So they continue to be subject to the demand for strong consistency.

Second, the very idea that some present-directed intentions escape the consistency demands to which most other intentions are subject seems to me not very plausible. After all, they are all equally intentions. Notice that we do not think belief works this way. That is, we do not see certain beliefs about the present as subject to weaker demands of consistency than beliefs about the future.

A second objection to (i) grants that there is a general presumption against such inconsistency, but urges that this presumption can sometimes be overridden and, indeed, is overridden in the present case. I have strong pragmatic reasons for intending to hit each target, since that is how I best pursue the reward. Given these pragmatic reasons to have both intentions, the fact that they fail to

be (given my beliefs) strongly consistent need not convict me of criticizable irrationality, contrary to (i).²¹

My response is to reject the contention that I must intend to hit each target in order best to pursue the reward. What I need to do is to *try* to hit each target. But this does not mean that I must *intend* to hit each target. Perhaps I must intend *something*—to try to hit each target, for example. But it seems that I can best pursue the reward without intending flat out to hit each target, and so without a failure of strong consistency.²² Given a presumption against such a failure, that is what I should do. If I nevertheless do intend to hit each target I am criticizably irrational. So (i) remains plausible.

What about (ii), the claim that what I do intentionally is *hit target 1*? Here the defender of the Simple View might urge that what I do intentionally is only to *hit one of the two targets*. So all that the Simple View requires is that I intend to hit one of the two targets. And that intention is not threatened by the demand for strong consistency.

In assessing this objection we must be careful to distinguish my case from other, superficially similar cases. For example, suppose there is a single target in front of you and you know it is either target 1 or target 2. But since the targets are labelled on the back you do not know which target it is. Still, you do know that you get a reward for hitting target 1 or for hitting target 2. So you shoot at, and hit the target in front of you, which turns out to be target 1.

Now, on one natural reading of 'trying', you were not trying specifically to hit target 1. You were only trying to hit whichever target it was that was in front of you. Further, on a natural reading of 'knowingly', you did not hit target 1 knowingly; for you did not know that it was target 1, rather than target 2, that you were hitting. Such observations make it plausible to say that while in hitting target 1 you intentionally hit one of the two targets, you did not intentionally hit target 1.

²¹I am indebted both to Kwong-loi Shun and to the editors of *The Philosophical Review* for forcing me to discuss this objection explicitly.

²²Perhaps there are other cases of trying to achieve each of two goals known to be incompatible in which, due to peculiarities of one's character, one really must intend to achieve each goal in order best to pursue each goal. But we need not suppose that my case is like this.

Again, suppose there are two targets close together, and one gun. You only have enough skill to aim in the vicinity of the pair of targets, trying to hit one or the other. And that is what you do. Suppose you hit target 1. Then it is plausible to say that in hitting target 1 you have intentionally hit one of the two targets without intentionally hitting target 1.

In both these cases, then, it might plausibly be insisted that you do not intentionally hit target 1. It is important to note, however, that my case is different from these. I am trying to hit each of two targets (though I am not trying to hit both). I am not just trying to hit a single target which, for all I know, is one or the other of two different targets. Nor am I just aiming the same shot at both targets in the attempt to hit one or the other. Rather, each of the two targets separately guides my attempt to hit it. Further, I know that if I successfully hit target 1 my endeavor to hit it will be terminated by my knowledge that I have hit that very target. So my case differs from yet a third variation in which I know, rather, that the machine will only tell me if one of the targets is hit, without telling me which one.²³ In this third variation it may be plausible to insist that all I do intentionally is hit one of the targets. But, again, my case is importantly different.

These contrasts with variant cases highlight features of my case which argue for the claim that I intentionally hit target 1. First, I want to hit target 1 and so am trying to do so. Second, my attempt to hit target 1 is guided specifically by my perception of that target, and not by my perception of other targets. Relevant adjustments in my behavior are dependent specifically on my perception of that target. Third, I actually hit target 1 in the way I was trying, and in a way that depends on my relevant skills. Fourth, it is my perception that I have hit target 1, and not merely my perception that I have hit a target, that terminates my attempt to hit it. Granted, if I had instead hit target 2 that also would have terminated my endeavor to hit target 1, given my knowledge of how the games are linked. Nevertheless, what actually does terminate my attempt to hit target 1 is my perception that I have hit that target. When all this is true it seems to me too weak just to say that I have intentionally hit one of the targets. Rather, I have intentionally hit target 1.

²³Example courtesy of the editors of *The Philosophical Review*.

Both crucial claims in my argument against the Simple View are, then, quite plausible. But this is not the end of the matter. We need also to know the larger theoretical advantages and disadvantages of giving up the Simple View. The remainder of this paper pursues some of these larger issues. Insofar as the alternative framework it sketches is independently plausible, it provides further support for the rejection of the Simple View.

3. INTENTION AND VOLITION

I have followed the Simple View in eschewing the reduction of intention to desire and belief. But I have also rejected the Simple View's detailed conception of the relation between intentional action and the state of intention. I now want to examine one natural way of arriving at the Simple View, a way that depends on a different sort of reduction—this time, of intention to volition. I want to do this for two reasons. First, such a reduction is a natural way of arriving at the Simple View, and so deserves some comment here. Second, this will allow me to show how one can be led to an alternative conception of what is common to both intention and intentional action. This alternative conception accepts the Single Phenomenon View, but supposes that the element common to both intentional action and the state of intention is volition, rather than intention itself. Since I will be defending a version of the Single Phenomenon View which sees intention as the common element (though, of course, not the version expressed in the Simple View), it is important for me to say why I reject this alternative.

When I *A* intentionally—and not merely by accident, by mistake, unwittingly, inadvertently and so on—it may seem plausible to say that I am in a sense “committed” to *A*-ing. There may seem here to be a kind of “practical commitment” to *A*-ing that goes beyond mere desire. This suggests that all cases of intentionally *A*-ing share a special pro-attitude in favor of *A*-ing, a pro-attitude distinct from an ordinary desire to *A*. The presence of this pro-attitude in favor of *A* guarantees the kind of commitment to *A*-ing characteristic of intentionally *A*-ing. We may call this special attitude *willing* or, alternatively, *volition*, and this suggestion the *Volitional Thesis*. On the Volitional Thesis, then, in intentionally *A*-ing I will to *A* (or, perhaps, that I *A*): I have a volition to *A* (that I *A*).

If one accepts the Volitional Thesis one needs to say more about willing. In particular, one needs to say what the relation is between willing to *A* and intending to *A*. Here we are faced with an important theoretical decision. On one conception willing and intending are completely distinct mental elements: my volition to *A* is itself neither an intention to *A* nor a necessary part of such an intention. When we see willing this way the Volitional Thesis, while compatible with the Simple View, provides no direct support for that view and could be accepted by one who rejected the Simple View. In this paper I leave open the question of the acceptability of the Volitional Thesis when willings are understood, in this way, as completely distinct from intentions.

More germane to present concerns is a second conception which supposes there to be a much tighter connection between intention and volition. One version of this second conception sees the volition to *A*, required by the Volitional Thesis for intentionally *A*-ing, as at least a necessary component of a present-directed intention to *A*. This is the *Necessity Thesis*. Finally, on an even stronger version of this conception a volition to *A*, in the sense of the Volitional Thesis, *just is* a present-directed intention to *A*. This is the *Identification Thesis*. And with the Identification Thesis we have arrived at the Simple View.

The Identification Thesis amounts to a reduction of present-directed intention to volition. Such a reduction seems fairly natural.²⁴ Yet, taken together with the Volitional Thesis it leads to the Simple View. Having rejected the Simple View we must block this reasoning at some point. Where?

²⁴Aune accepts the Identification Thesis in *Reason and Action*, op. cit., Chapter II, Section 4. Searle also seems to be guided by some such reduction in *Intentionality*, op. cit. His initial arguments for the presence of an "intention in action" in all intentional action are just the arguments commonly used to argue for the presence of volitions, for example: James's case of the anesthetized patient who mistakenly thinks that he is raising his arm (p. 89). Such cases suggest that intentional action involves volitions in roughly the sense of the Volitional Thesis. By labelling this volitional element "intention in action" Searle takes the *further* (and, so far as I can see, unargued) step of identifying it with present-directed *intention*. Finally, Castañeda explicitly says that present-directed intentions are volitions (*Thinking and Doing*, op. cit., p. 277), though in this book he accepts only the weaker view (described above in note 4) of what present-directed intentions (i.e., volitions) are required in intentional action.

Return to the last video games example. This example does not threaten the Volitional Thesis taken by itself. It remains open, for all that that thesis says, that I both will to hit target 1 and will to hit target 2. This is because by itself the Volitional Thesis offers no reason for supposing that willings are subject to the same demands of strong consistency to which intentions are subject. Nor does the example directly challenge the idea that willings, of the sort required by the Volitional Thesis, are necessary components of corresponding intentions. What the example precludes is that willings of the sort required by the Volitional Thesis be *identified with* corresponding intentions. If my willings to hit each target were just present-directed intentions to hit them, I would be criticizably irrational; but I am not. To avoid the Simple View we must reject the reduction of present-directed intentions to volitions of the sort required by the Volitional Thesis.

In light of our discussion, we can see what goes wrong with such a reduction. The Volitional Thesis introduces the notion of volition to capture the special commitment it supposes to be characteristic of intentional action. In contrast, the idea of an intention to act is partly tied to future-directed intentions and plans, and to their characteristic commitment to future action. When we identify present-directed intentions with such volitions we implicitly assume that these two roles do not bring with them conflicting demands. But what we learn from the video games example is that they do. So we should reject this reduction.

Having rejected the Identification Thesis, could we still retain the Necessity Thesis, the view that a volition to *A* is a necessary part of a present-directed intention to *A*? This would be to see a volition to *A*, rather than a full-blown intention to *A*, as the element common to both intentionally *A*-ing and having a present-directed intention to *A*. This would lead naturally to an alternative version of the Single Phenomenon View. On this alternative version the basic, single phenomenon is volition; it is a volition to *A* that is common to both my intentionally *A*-ing and the intention to *A*. While I may intend to *A* in intentionally *A*-ing, I need not: I need only will to *A*. An intention to *A* is a volition to *A* *together with something else*.

What else? In his important British Academy lecture²⁵ H. P.

²⁵Op. cit.

Grice in effect pursues a version of this strategy. Grice first introduces a general notion of willing which has the feature that I will that I *A* whenever I either intentionally *A* or intend now to *A* later. Gricean willings, while embodying (in the present-directed case) the special commitment characteristic of intentional action, are not confined to the present. Grice then goes on to claim that my intention to *A* is my willing that I *A* together with my belief that I will, as a result, *A*. There is a single phenomenon involved in both intentional action and the state of intention, but it is not intention itself. It is, rather, volition—understood as a proper part of intention.

This view has several virtues. It avoids the identification of present-directed intention and volition that we have seen to founder on the demand for strong consistency of intention. At the same time it provides an explanation of why intentions are subject to such consistency demands, namely: because beliefs are, and an intention to *A* includes the belief that one will. In this way it provides for a more complex connection between the commitment characteristic of intentional action, and that characteristic of future-directed intentions and plans, than is allowed by the Simple View.

The problem is that this view requires a return to the strong belief requirement on intention. And we have seen reason to be doubtful of that requirement. Further, there is no obvious way to weaken this belief requirement without creating other difficulties. For example, suppose we try saying that an intention to *A* is a volition that one *A* together with a belief that, as a result, one is more likely than not to *A*. The problem now will be that we have undermined the general capacity of rational intentions to be unified into larger, rational plans. This is because we have now blocked the inference from

- (a) I rationally intend to *A* and rationally intend to *B*.

to

- (b) It would not be irrational for me to intend to *A* and *B*.

It is not generally true that if I rationally believe of *each* of *A* and *B* that I am more likely than not so to act, I can rationally believe the same of my performing *both* actions. So the inference from (a) to (b)

will not go through. But, as we have seen, intentions are at least potentially elements in larger coordinating plans. To be rational my intentions should at least be capable of combining into a larger plan that is not irrational. So we will want to retain the inference from (a) to (b).

In light of these difficulties perhaps we should put aside such belief requirements on intention and appeal directly to the constraint of strong consistency. One way to do this might be to say that my intention to *A* is my volition to *A* together with my disposition to impose this constraint on that volition.

The problem now is that we make the step from volition to intention appear arbitrary. Recall that it is not generally appropriate to impose the demand for strong consistency on one's volitions: this is the lesson of the video games example. So, on the suggested view, to intend to *A* I must be disposed to treat my volition to *A* in a special way, a way not generally appropriate for volitions. But, having given up the belief condition on intention, the suggested view leaves us with no explanation of why this special treatment is appropriate in this case. We have no explanation of why I should impose the demand for strong consistency on this volition but not on others.

The strategy of constructing intentions out of volitions and other things, we now see, faces a dilemma in providing for the consistency constraints characteristic of intention. If it tries to account for these constraints by adding to volition a further belief condition, this condition will be too strong. If it tries just to tack onto volition a disposition to impose the relevant consistency constraints, it makes the step from volition to intention seem arbitrary. Faced with this dilemma, I propose taking a different tack.

4. INTENTION AND MOTIVATIONAL POTENTIAL

Both the Simple View and the Volitional Thesis agree in supposing that—*pace* the desire-belief model—intentional action involves a special pro-attitude, distinct from the agent's desires and beliefs. The problem has been to say more precisely what that special attitude is. On the Simple View it is an intention to act in the way one acts intentionally; and we have seen this idea to be at odds with the requirement that intentions be strongly consistent. On the Voli-

tional Thesis all intentional action involves a distinctive volition so to act. I have not criticized this view directly. I have instead focused my critical attention on attempts to weld such a view to either a reduction of present-directed intention to volition (which just returns us to the Simple View) or the Gricean view of volitions as proper parts of intentions.

Both the Simple View and the Volitional Thesis share a common assumption. They both assume that if there is a distinctive pro-attitude involved in intentionally *A*-ing, it will be a pro-attitude specifically in favor of *A*—that there must be a tight fit between what is done intentionally and what is intended (willed). This is the *assumption of tight fit*. Together with our video games example, this assumption leads us to reject the idea that intentional action generally involves an intention, that intention is the element common to both the state of intention and intentional action.

I propose to give up the assumption of tight fit and to distinguish between what is intended, and the sorts of intentional activity in which an intention may issue. Making this distinction, we can say that when I *A* intentionally I intend *something*, but I may not specifically intend *to A*. Our notion of intentional action embodies a complex scheme for the classification of actions (or, perhaps, actions “under a description”). To understand the relation between intention and intentional action we must recognize that the factors that determine what is intended do not completely coincide with the factors that, on this scheme, determine what is done intentionally.

Recognizing this, we can accept a version of the Single Phenomenon View which sees intention as the common element in both intentional action and the state of intention. To find a common element we need not retreat to some proper part of intention, volition. Actions are intentional in part because of their relations to intentions. But the admissible relations are more complex than those envisaged by the Simple View.

In the theory of action one can be led into two different mistakes (among others!). The first, built into the desire-belief model, is to suppose that intentional action involves no distinctive state of intention at all. The second, made by the Simple View, is to suppose that intentional action always involves an intention *so to act*—a supposition that does not do justice to the role of intentions in coordinating plans. I am proposing a way between. In acting inten-

tionally there is something I intend to do; but this need not be what I do intentionally.²⁶

Supposing, then, that there are cases in which I intentionally *A* and yet do not intend to *A* but only intend to *B*, for some appropriate *B*, a full account of our scheme for classifying actions as intentional will need to sort out just when this can be so. Whatever its details, such an account will implicitly specify a four-place relation between intentions, desires, beliefs and types of actions. It will say what types of actions may be performed intentionally in the course of executing a certain intention, given a certain background of desires and beliefs. This allows us to define a useful notion, that of the *motivational potential* of an intention. *A* is in the motivational potential of my intention to *B*, given my desires and beliefs, just in case it is possible for me intentionally to *A* in the course of executing my intention to *B*. If I actually intend to *A* then *A* will be in the motivational potential of my intention. But we need not suppose

²⁶Are there cases of spontaneous activity that, while plausibly classified as intentional, do not involve anything reasonably identifiable as an intention to act? If you unexpectedly throw a ball to me I might reach up and catch it. I catch it intentionally, but perhaps my catching it involves no intention to do something.

My worry here is not based on the false assumption that all present-directed intentions are preceded by corresponding, future-directed intentions. My worry, rather, is that once we see what a present-directed intention is (in part by reflecting on future-directed intentions) it may not be obvious that all spontaneous action that is intentional (and not mere reflex behavior, as when I blink at the oncoming ball) must involve such a state. Perhaps our scheme for classifying actions as intentional, while treating as central actions involving intentions, is more inclusive than that. If this were so then we would have to limit my version of the Single Phenomenon View to those central cases of intentional action. Such a limitation on the Single Phenomenon View would still be compatible with (though it would not require) the claim of the Volitional Thesis that what distinguishes my intentional catching of the ball from my blinking is the role played by an appropriate volition.

It is not possible to address this worry in a definitive way without a more detailed specification of the relations between intention and action that can make that action intentional, without a full account of what I will be calling *motivational potential*. For example, lacking such an account it is unclear whether we can appeal to a general intention to protect myself from flying objects to explain, compatibly with the present account, why my catching the ball is intentional. Since I do not offer such a full account of motivational potential here, I leave the resolution of this matter to another occasion.

that if *A* is in the motivational potential of an intention of mine then I intend to *A*.

Consider the last video games example. My intention includes my hitting target 1 in its motivational potential: it is possible, given my desires and beliefs, for me to hit target 1 intentionally in the course of executing my intention. Nevertheless, I do not intend flat-out to hit target 1. While hitting target 1 is in the motivational potential of my intention, it is not *what* I intend.

What then do I intend? There are several possibilities. I might intend to try to hit target 1, and also to try to hit target 2. I might intend to hit target 1 if I can, and similarly concerning target 2. I might even just intend to hit one of the two targets; though we must be careful to distinguish this case from the cases discussed in Section 2 in which, though I intend to hit one of the two targets, my intention does not include hitting target 1 in its motivational potential. The important point is just that my intention may include hitting target 1 in its motivational potential without including it in what is intended.

That my intention includes hitting target 1 in its motivational potential, even though it is not an intention to hit target 1, does not by itself *explain* why it is true that I hit target 1 intentionally. This is clear from the definition of motivational potential. The notion of motivational potential is intended to *mark* the fact that my intention to *B* may issue in my intentionally *A*-ing, not to explain it. It is a *theoretical placeholder*: it allows us to retain theoretical room for a more complex account of the relation between intention and intentional action while leaving unsettled the details of such an account. Such an account would not itself use the notion of motivational potential but would, rather, replace it with detailed specifications of various sufficient conditions for intentional conduct.

Let me put the point this way. On the theory just sketched, if I *A* intentionally then I *A* in the course of executing some intention to *B* and, given my desires and beliefs, this intention contains *A* in its motivational potential. This means that there will be some true statement(s) along the lines of:

If *S* intends to *B* and *S* *A*'s in the course of executing his intention to *B* and _____, then *S* *A*'s intentionally.

A full-blown theory of intentional action will tell us how such blanks should be filled in. For example, our discussion of the video games example suggests that one such specification of sufficient conditions would be roughly along the lines of the following:

S intentionally *A*'s if

- (1) *S* wants to *A* and for that reason intends to try to *A*, and
- (2) *S* *A*'s in the course of executing his intention to try to *A*, and
- (3) *S* *A*'s in the way he was trying to *A*, and
- (4) (2) and (3) depend, in an appropriate way, on *S*'s relevant skills.

Without working out the details, we can see that such a specification would use conditions like (3) and (4) to fill in the theoretical space opened up by our distinction between what is intended and what is in the motivational potential of an intention.

This new theoretical space allows us to formulate a more satisfactory alternative to the desire-belief model than those so far considered. In contrast with the desire-belief model, we can grant that intentional action at least typically involves a distinctive pro-attitude that is not reducible to the agent's desires and beliefs. In particular, intention is a distinctive pro-attitude involved in intentionally *A*-ing, though it need not be an intention *to A*. By allowing this flexibility in what is intended we do better than the Simple View in providing for the consistency demands on intentions. We can allow, for example, that when I intentionally hit target I what I intend need not involve me in inconsistency.

This flexibility also takes away a main source of motivation for accepting the Necessity Thesis and treating an intention to *A* as consisting of a volition to *A* plus something else. Having given up the assumption of tight fit, we no longer must choose between an intention to *A* and a volition to *A*—understood as a proper part of such an intention—in order to locate a distinctive pro-attitude generally involved both in an intention to *A* and in intentionally *A*-ing. Further, since all intentions are subject to a demand for strong consistency, we avoid an analogue of the puzzle, faced by the defender of the Necessity Thesis, about why we should impose such constraints on only some proper sub-set of our volitions.

In response one might still worry that the distinction, between what is intended and what is in the motivational potential of an

intention, is illusory. As Anscombe famously remarks, “the primitive sign of wanting is *trying to get*.”²⁷ But what is true about wanting seems even more clearly true about intention: the “primitive sign” of an intention to *A* is trying to *A*. In the face of this I have tried to drive a wedge between an intention whose execution may involve both trying to *A* and intentionally *A*-ing, and an intention to *A*. I have claimed that one might have the former intention and yet still not intend to *A*. But how is that possible? Differences in what I intend should reveal themselves in differences in the roles played by my intentions. But the basic role present-directed intentions play is in motivating and guiding present conduct. So it may seem unclear that there is a real difference between intending to *A* and having an intention whose role includes the motivation of intentionally *A*-ing.

The response to this worry is that intentions play other important roles. Differences in these roles can discriminate between two intentions, both of which include *A* in their motivational potential but only one of which is an intention to *A*. That there are these other important roles is clear from the methodological priority of future-directed intention; for a basic role played by future-directed intentions is as elements in coordinating plans. There are differences in the role played in such plans by an intention to *A* and that played by other intentions which include *A* in their motivational potential. Included among these will be differences in the constraints imposed on yet other intentions, given the demand for strong consistency. What I intend, when I have a future-directed intention, will be in part reflected in the ways in which my intention constrains my other intentions by way of this consistency demand. Thus, if my future-directed intentions concerning targets 1 and 2 do not convict me of criticizable inconsistency then, given my beliefs, they are not intentions to hit target 1 and to hit target 2. This is so even though my intention concerning target 1 includes hitting it in its motivational potential, and similarly with my intention concerning target 2.

A similar point applies to present-directed intentions. What I intend when I have a present-directed intention will not be simply a matter of the sorts of intentional conduct in which my intention

²⁷Op. cit., p. 68.

might issue. I can have a present-directed intention which includes hitting target 1 in its motivational potential even though I do not intend flat-out to hit target 1. For my intention to be an intention *to hit target 1* it must constrain my other intentions accordingly, by way of the demand for strong consistency. And, as we have seen, my intentions concerning targets 1 and 2 may have hitting each target in their motivational potential without constraining each other in the ways characteristic of intentions to hit these targets.

5. MOTIVATIONAL POTENTIAL EXTENDED

Let us sum up so far. Desires, beliefs and intentions are basic elements in the commonsense psychology underlying intentional action. Intentions are typically elements in plans. Intentional action generally involves an intention to act. The state of intention is itself the common element in both the states and the actions included within our conception of intention: the Single Phenomenon View is correct. The intention involved in intentional action need not, however, be an intention so to act. My intention may include *A* in its motivational potential even though I do not, strictly speaking, intend to *A*. The coherence of this latter idea is ensured by the role intentions play in coordinating plans. All this is neutral on the question of whether intentional action involves a special volitional element that is completely distinct from intention. But it does eliminate the need to introduce volitions as special psychological elements related to intentions as part to whole, and serving as the common element in intentional action and intending to act.

This approach depends on driving a wedge between what I intend and the motivational potential of my intention. Now, the wedge I have so far argued for has been rather thin: it has directly concerned only certain special cases in which the demand for strong consistency created problems for the Simple View. But once we have this wedge we can widen it in ways that promise to be useful. Let me briefly sketch two such ways.

Suppose I intend to run the marathon and believe that I will thereby wear down my sneakers. Now it seems to me that it does not follow that I intend to wear down my sneakers, and in a normal case I will not so intend. One sign of the absence of such an intention will be the fact that I am not at all disposed to engage in

further reasoning aimed at settling on some means to wearing down my sneakers. In contrast, if I intended to get to the track by 9 a.m., as a means to running the race, I would be disposed to engage in reasoning aimed at figuring out how to do that.²⁸ My attitude towards wearing down my sneakers does not play the role in further means-end reasoning that an intention to wear them down would normally play.

Even so, if I proceed to run the marathon and actually do wear down my sneakers then I might well do so intentionally. Perhaps this is clearest in a case with two further features.²⁹ First, I not only believe I will wear them down; I consciously note this while I am running. Second, wearing them down has some independent significance to me; perhaps they are a family heirloom. In a case with these two further features I think we would classify my action as intentional. Yet it does not seem that these further features must change *what I intend* in running the race. Given my relevant beliefs and desires, in executing my intention to run the race I may intentionally wear down my sneakers; and this even though I do not intend to wear them down. So while *what I intend* does not include wearing down my sneakers, the motivational potential of my intention does.

Generalizing, we can expect a full theory of intentional action to generate true statements along the lines of

If *S* intentionally *B*'s in the course of executing his intention to *B*, and *S* believes that his *B*-ing will result in *X*, and his *B*-ing does result in *X* and _____, then *S* intentionally brings about *X*.

²⁸I introduce this further intention to make it clear that I am not just denying that I intend to wear down my sneakers "as an end." I do not intend to get to the track by 9 a.m. as an end; but I still do intend to do so. In contrast, I may not intend at all to wear down my sneakers. For probing discussions of related matters see Jonathan Bennett, "Morality and Consequences," *The Tanner Lectures on Human Values*, edited by Sterling M. McMurrin (Cambridge: Cambridge University Press, 1980), lecture III; and Gilbert Harman, "Rational Action and the Extent of Intentions," *op. cit.*

²⁹As Allan Gibbard helped me see. Note that I do not say that I run the race *with the intention* of wearing down my sneakers. I do not discuss acting with a further intention in this paper.

For present purposes we can leave aside the subtle issue of just how the blank should be filled in (e.g., must it add that *S* is aware that he is bringing about *X* and is not indifferent as to whether or not he does bring it about?). The important point is that these sufficient conditions will not include the requirement that *S* actually intends to bring about *X*. This means that motivational potential can be extended by our beliefs about the upshots of what we intend, even when what we intend is not thereby extended.

Consider a second sort of case. I intend to shoot a jump shot. I know that my jump shot will have to contain certain sub-components, for example: stopping on my left foot. But as a skilled jump-shooter I need not intend all this, for my intentions and plans are typically at a level of abstraction appropriate to my skills. I may just intend to shoot the jump shot, perhaps as part of a larger plan to score and then to try to steal the in-bounds pass.

We may say that my stopping on my left foot is a *necessary constitutive means* of my shooting the jump shot. What this case suggests is that I may, while guilty of no criticizable irrationality, intend to *B*, know that *A* is a necessary constitutive means of *B*-ing, and yet not intend to *A*. Rational intention need not be transmitted along the lines of known, necessary constitutive means.

Nevertheless, it seems that the motivational potential of my intention may be transmitted along such lines even when what I intend is not. If I successfully execute my intention and shoot the jump shot, and if in so doing I stop on my left foot, then I may well have stopped on that foot intentionally. So the motivational potential of my intention to shoot the jump shot may include stopping on my left foot.

Here again this may be clearest for cases which have two further features. First, I not only know I must stop on my left foot; I consciously note this as I am shooting. Second, stopping on my left foot has some independent importance to me; perhaps I have recently injured it and it behooves me to go easy on it. In a case with these two further features I think we would classify my stopping on my left foot as intentional. Yet it does not seem that these further features force a change in what I intend. What I intend may remain just to shoot the jump shot. But given my background of beliefs and desires my intention includes stopping on my left foot

in its motivational potential: it is possible for me to stop on my left foot intentionally in the course of executing my intention.

Generalizing again, we can expect a theory of intentional action to issue in true statements along the lines of

If *S* intentionally *B*'s in the course of executing his intention to *B*, and *S* believes that his *A*-ing is a necessary constitutive means of his *B*-ing and *S* *A*'s in the course of executing his intention to *B* and _____, then *S* intentionally *A*'s.

Here again the important point is not the details concerning how to fill in the blank, but just that an intention to *A* is *not* required. This means that motivational potential can be extended by means-end beliefs, even when what is intended is not thereby extended.

These cases illustrate some of the complexities of our scheme for the classification of actions as intentional. The Simple View forces us to read these complexities back into the agent's intentions: it includes in what is intended everything done intentionally. Our view loosens the connection between what is intended and what is done intentionally: it sees what is intended as a fact about the agent's mind which need not reflect all the complexities of our scheme for classifying actions as intentional. It does this by using the notion of motivational potential to provide a buffer between the considerations that influence the intentionality of action and those that influence what a person intends.

6. MOTIVATIONAL POTENTIAL AND THE DISTINCTIVENESS OF INTENTION

I now want to argue that this buffer helps support the central claim that intentions are distinctive states of mind. It does this by protecting regularities important to the defense of this claim.

The classificatory schemes involved in our commonsense framework play certain roles in our lives, and we can expect the details of such schemes to be shaped by those roles. An important role played by our scheme for classifying actions as intentional is that of identifying ways of acting for which an agent may be held responsible: our concern is not limited to the description and explanation of actions, but extends to the assessment of agents. This is why it seems natural to classify as intentional my wearing down my

sneakers. After all, as Sidgwick notes in defending his proposal to “include under the term ‘intention’ all the consequences of an act that are foreseen as certain or probable”: “we cannot evade responsibility for any foreseen bad consequences of our acts by the plea that we felt no desire for them.”³⁰

Now, the case for seeing intentions as distinctive states of mind depends on locating them in an explanatory system connecting environment and behavior, and on identifying their distinctive role in this system. To do this there need to be underlying regularities connecting intentions with each other and with other states and processes.³¹ Further, these regularities must be significantly dependent on *what* is intended; a regular connection between, say, intentions formed during winter quarter and nervousness is not the sort of regularity we need. To the extent to which our scheme for determining what is intended is shaped by our concern, not only with explanation of action, but with the assignment of responsibility, it will be harder to find such regularities. This is because such a concern would tend to lead to the ascription of intentions which do not play their normal roles in motivation and practical reasoning.

To see this, consider again my intention to get to the track by 9 a.m., as a means to running the race. This intention plays a pair of roles important to attempts at explanation. First, it triggers further means-end reasoning concerning how to get to the track by then. Second, when the time comes it motivates activity guided by my beliefs (many of them perceptual) about where the track is.³² In these respects it contrasts with my mere expectation that I will wear

³⁰*The Methods of Ethics* (seventh edition) (New York: Dover, 1966), p. 202. For a useful discussion of such matters, and their relations to some views of Harman’s, see Neil Lubow, “Acting Intentionally,” unpublished manuscript.

³¹It is not to my purpose here to discuss how strict these regularities need be, but I would expect them only to be of a sort involved in what Grice calls “*ceteris paribus* laws.” See H. P. Grice, “Method in Philosophical Psychology (From the Banal to the Bizarre),” *Proceedings and Addresses of the American Philosophical Association* 48 (1974-75), pp. 23–53.

³²Of course, to play this motivational role my intention need not be an intention to get to the track by 9 a.m. It might just be an intention to try, or to get there by then if my old car holds up. Still, if I do intend to get there by 9 a.m. (and do not merely expect that I will) my intention will normally play the cited role.

down my sneakers as a result of running. I am neither disposed to engage in reasoning aimed at settling on a means to wearing them down, nor do I guide my running of the race by keeping track of the state of my sneakers.³³

There are, then, distinctive regularities connecting what is intended with further practical reasoning and with what beliefs guide our activity. The Simple View undermines such regularities. By reading back from the intentionality of my wearing down my sneakers to an intention to wear them down, it ascribes to me an intention which is outside the web of these regularities; for my attitude toward wearing down my sneakers does not play the roles characteristic of an intention to do so. To support such regularities we need to allow our concern with responsibility to shape what is done intentionally without similarly shaping what is intended. We need to allow our concern with responsibility to lead us to classify my wearing down my sneakers as intentional, without forcing us to say that I intend to wear them down. This is what the notion of motivational potential allows our theory to do.

Returning to our video games example, we can make a similar point. Here the relevant regularity is a general tendency towards equilibrium. Generally, when an agent notices that his intentions fail to be strongly consistent there will be an attempt at revision, aimed at achieving consistency. But this regularity is undermined if we suppose that in cases such as our video games example there are strongly inconsistent intentions and yet no tendency towards appropriate revision. The notion of motivational potential allows us to protect this regularity and yet still grant that I hit target I intentionally.

7. TWO FACES OF INTENTION

Intention is Janus-faced, tied both to intentional action and coordinating plans. I have tried to sketch a version of the Single Phenomenon View that provides room for both of these faces of

³³I *might* guide my running of the race by keeping track of the state of my sneakers—for example, if I use them as a pedometer. Since even then I would not intend to wear them down, the presence of such guidance does not ensure intention. My point here is only that its absence indicates an absence of intention.

intention, and for an appropriate link between them. In doing this I have tried to avoid the oversimplifications of the Simple View and the Identification Thesis, as well as the difficulties that arise when we try to construct intentions out of volitions and other things. I have also tried to leave room for the different effects which our concern with the ascription of responsibility has on the different classificatory schemes included within our conception of intention. And, finally, I have tried to do this in a way that recognizes, exploits and supports the distinctiveness of an agent's intentions and plans.³⁴

Stanford University

³⁴An ancestor of this paper was read at a conference on practical reasoning held at The University of Dayton in March, 1983. I benefitted from the careful remarks of my commentator, Donald Gustafson. Other versions have been read at colloquia at Stanford University and at San Jose State University. John Perry's detailed comments on a late version of this paper helped me to clarify several important points, as did comments from the editors of *The Philosophical Review*. Among the many other people who helped me with this paper I want particularly to thank Gerald Barnes, Arnold Davidson, John Dupre, Allan Gibbard, Gilbert Harman, David Hilbert, David O'Conner, Denis Phillips, Adrian Piper, Kwong-loi Shun, and Howard Wettstein. Some of the initial work on this paper was done with the support of a NEH summer research grant.