

controlling present action and to the roles of future-directed intention in the interim between intention-acquisition and intention-execution. By unpacking this complex of dispositions and norms we demystify commitment without allowing it to collapse into mere desire and expectation. We explain it without supposing that it involves a ghostly mode of influence of future conduct by present decisions, without requiring that a commitment to action be irrevocable, and without making it a mystery why we ever bother to form intentions for the future.

Now, some philosophical work in the theory of action reveals a tendency to see intention as either a special kind of evaluation or a special kind of belief. Donald Davidson's later work on intention is an example of the former sort of view; Stuart Hampshire and H. L. A. Hart's is an example of the latter.³ On the view of intention sketched here intention is a state characterized by a complex of functional roles and associated norms. These roles and their associated norms constitute the special commitment characteristic of intention. Such roles involve, of course, systematic relations with attitudes of evaluation and belief. But on the present theory we need not try to assimilate intention to either kind of attitude. We can just say that an intention to act is a complex form of commitment to action, a commitment revealed in reasoning as well as in action. To say this we must take seriously dispositions to reason in certain ways, and the functional roles of attitudes both in reasoning and in other processes that are part of the psychological background against which practical reasoning takes place. And this highlights once again the importance of practical reasoning and planning to our understanding of intention.

chapter 8

Two Faces of Intention

I HAVE been sketching a planning conception of intention. On this conception a central fact about us is that we are planning creatures. We frequently settle in advance on plans for the future. 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: as such, intentions are distinctive states of mind, not to be reduced to clusters of desires and beliefs.

Now it is time to return to the relation between intention and intentional action. Recall the observation with which I began this book. 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.

8.1 The Simple View

How should the planning theory of intention view this relation? I may intend to start my car today: this is a future-directed intention. But I may

also intend to start my car beginning 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 to start it, it seems that what I intend will include starting it.

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*. The Simple View does not say that there must be a separate event of intending to *A* for each intentional *A*-ing. But it does require that one's mental states at the time one intentionally *A*'s include an intention to *A*.¹

Let me clarify this further. Suppose I intentionally start my car. On the Simple View it follows that

- (1) I intend to start my car.

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

- (2) *aRb*

where *b* is replaced by a singular term denoting an actual, particular action of my starting my car.²

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 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 toward which common sense initially leans, as well as a view implicit in many discussions of intention in moral philosophy.³ Nevertheless, although 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 8.2 and 8.3 I explain why. In Sections 8.4 through 8.7 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 continues to see intentions as distinctive, and to see 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 although 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 this chapter by arguing that this distinction has a further important 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.

8.2 The Argument against the Simple View

My argument against the Simple View is rooted in my conception of intentions as elements in larger, coordinating plans, and the resulting demands for strong consistency. Recall that to be well suited as elements of such coordinating plans my intentions need to be (other things equal) strongly consistent, relative to my beliefs. It should be possible (other things equal) for me successfully to execute all my intentions in a world in which my beliefs are true. But once we take this demand for consistency seriously, the Simple View is subject to challenge.

The challenge can be presented in terms of a series of three examples.⁴ 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 skill 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.⁵

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 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.⁶ 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" to 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 already noted, 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 my 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 I hit one of the targets, both games are over. If both targets are about to be hit simultaneously, the machines just shut down and I hit neither target. Both targets remain visible to me; so I can see which target I hit if I hit either one. And there is a reward for hitting either target. But I know that although 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 skill 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. Yet it seems clear that 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 irra-

tionality, I do not have both of these intentions. Since my relevant intention in favor of hitting target 1 is the same as that 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.⁷

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, one that I will discuss briefly.

Suppose I intend now to go to the concert tonight. What must I believe about my future concert-going? As I noted in Chapter 3, some philosophers accept the strong thesis that I must now believe I will go.⁸ Once we are given this strong belief requirement on future-directed intention, however, it is 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.⁹

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 about 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.¹⁰ 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. For reasons I indicated in Chapter 3, however, it seems to me that this strong belief condition is problematic in ways in which the demand for strong consistency is not. Though there is notorious disagreement on this matter, it seems to me plausible to

suppose that sometimes intentions just do not satisfy such a strong belief condition. At the least 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*.

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* in order 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 act in a certain way while 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 so 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.

8.3 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:

- (a) If in this case I had present-directed intentions which failed to be strongly consistent, I would be criticizably irrational.
- (b) 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 (a). Such a defender might urge that, for all I have said, only *future-directed* intentions are subject to the strong consistency requirement. So, contrary to (a), I can intend to hit target 1 *now*, and similarly target 2, without being criticizably irrational.

This objection is inadequate for two reasons. First, the argument for the demand for strong consistency depends on the observation that intentions typically play a coordinating role as elements of larger plans. Now, although this role is clearer in the case of future-directed intentions, it is also important in connection with present-directed intentions. Sup-

pose 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 the games. When the time comes, these become present-directed intentions. But they continue to be part of my coordinating plan, and so they continue to be subject to the demand for strong consistency.

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

A second objection to (a) 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 strongly consistent (given my beliefs) need not convict me of criticizable irrationality, contrary to (a).¹¹

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 shoot at 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.¹² 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 (a) remains plausible.

What about (b), the claim that what I do intentionally is to *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 labeled 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 although in hitting target 1 you intentionally hit one of the two targets, you did not intentionally hit target 1.

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.¹³ In this third variation it may be plausible to insist that all I do intentionally is to 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.

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. In the remainder of this chapter and in Chapter 9 I will

pursue some of these larger issues. Insofar as the alternative framework I sketch is independently plausible, we have further support for the rejection of the Simple View.¹⁴

8.4 Intention and Motivational Potential

The Simple View supposes that there must be a tight fit between what is done intentionally and what is intended. I propose to give up this assumption 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, one which sees intention as the common element in both intentional action and the state of intention. Actions are intentional in part because of their relation 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 these two. In acting intentionally there is something I intend to do; but this need not be what I do intentionally.

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 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. Though 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 shoot at target 1 and also to shoot at target 2. 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 in the case of target 2. I might even just intend to hit one of the two targets; but we must be careful to distinguish this case from the cases discussed in Section 8.3 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, rather than 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 would tell us how such blanks should be filled in. For example, our discussion of the video-games ex-

ample suggests that one such specification of sufficient conditions would be roughly along the following lines:

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) conditions (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 the 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, on the one hand, and the Simple View, on the other. In contrast with the desire-belief model, we can grant that intentional action involves a distinctive pro-attitude—namely, intention—that is not reducible to the agent's desires and beliefs. But, in contrast with the Simple View, the intention that is involved in intentionally *A*-ing 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 1 what I intend need not involve me in inconsistency.

One might object that this distinction, the distinction between what is intended and what is in the motivational potential of an intention, is illusory. As Anscombe remarks, "The primitive sign of wanting is *trying to get*."¹⁵ 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, one might object, 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 that 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, roles I have been emphasizing throughout this book. Differences

in these roles can serve to discriminate between two intentions, both of which include *A* in their motivational potential but only one of which is an intention to *A*. Future-directed intentions are elements in coordinating plans and inputs into further practical reasoning. In terms of the role played in such plans and reasoning, an intention to *A* differs from other intentions which include *A* in their motivational potential. Among these differences are 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 through 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 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 include hitting each target in their motivational potential without constraining each other in the ways characteristic of intentions to hit these targets.

8.5 Motivational Potential Extended

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 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 last idea is ensured by the role intentions play in coordinating plans.

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 a way that promises to be useful.

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 9:00 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.¹⁶ My attitude toward 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 including two further features.¹⁷ 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 containing 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 although *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*.

For present purposes we can leave aside the subtle issue of just how the blank should be filled in (for example, 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 intend 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.¹⁸

I will be returning to related issues about expected but nonintended upshots of intended actions in Chapter 10. For now I want to emphasize 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. My version of the Single Phenomenon 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.¹⁹ 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.²⁰

8.6 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 and to the appropriateness of reactive attitudes like indignation and resentment.²¹ 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."²²

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. 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. 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 9:00 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.²³ In these respects it contrasts with my mere expectation that I will wear 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.²⁴

There are, then, distinctive regularities connecting what is intended with further practical reasoning and with the beliefs that 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 my theory to do.

Returning to my video-games example I can make a similar point. Here the relevant regularity is a general tendency toward equilibrium. Generally, when an agent notices that his intentions fail to be strongly consistent, he will make an attempt at revision, an attempt aimed at achieving consistency. But this regularity is undermined if we suppose that in cases such as the video-games example there are strongly inconsistent intentions and yet no tendency toward appropriate revision. The notion of motivational potential allows us to protect this regularity and yet still grant that I hit target 1 intentionally.

The point is not that our scheme for classifying actions as intentional is any less important than our scheme for characterizing states of mind as intentions. Nor am I claiming that the former is somehow gratuitous

in a way in which the latter is not.²² My point is rather that these schemes serve different interests and cut different ways. When we attempt to connect these schemes in an overly simple way—as the Simple View does—we do violence to this complexity.

8.7 Spontaneous Intentional Action

I now want briefly to consider one final objection to my version of the Single Phenomenon View. Suppose you unexpectedly throw a ball to me and I spontaneously reach up and catch it. On the one hand, it may seem that I catch it intentionally; after all, my behavior is under my control and is not mere reflex behavior, as when I blink at the oncoming ball. On the other hand, it may seem that, given how automatic and unreflective my action is, I may well not have any present-directed intention that I am executing in catching the ball. The worry here is not based on the false assumption that all present-directed intentions are preceded by corresponding, future-directed intentions. The worry, rather, is that once we see what a present-directed intention is (in part by reflecting on future-directed intentions and plans), it is unclear that such unreflective, spontaneous behavior 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.²⁶

In response two points need to be made. First, many such spontaneous actions might best be characterized as actions that, although they are under the agent's voluntary control and are purposive, still are not intentional (even though they are not unintentional). Second, it may be in some cases that in catching the ball I am executing some long-standing personal policy—for example, a policy of protecting myself from flying objects. I do not have a present-directed intention specifically to catch this very ball, but my action still involves an intention, namely: my general intention to protect myself in such circumstances.²⁷

The first point allows us to block at least some potential counterexamples; for such cases of purposive and voluntary—but still not intentional—behavior do not fall under my version of the Single Phenomenon View. The second point suggests that the intention involved in spontaneous intentional action need not itself be present-directed, but may rather be a long-standing general intention, one that applies to the present case.²⁸ Taken together these points go at least part of the way toward answering the present objection. But matters here are complex, and I am unsure whether such a defense can work for all cases. Perhaps there will remain cases of spontaneous intentional action that fall outside my ver-

sion of the Single Phenomenon View. Still, it does seem that the present theory will apply to the vast majority of the cases of intentional activity of planning agents like us. And that is good enough for my purposes.

8.8 Janus-faced Intention

Let me sum up. Intention is Janus-faced, tied both to intentional action and coordinating plans. In this chapter I have sketched a version of the Single Phenomenon View that provides room for both of these faces of intention and for an appropriate link between them. In doing this I have tried to avoid the oversimplifications of the Simple View. 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.

Throughout this discussion I have focused on the use of the adverb “intentionally” to characterize action. But we also characterize actions by speaking of an intention *with which* the agent acted. Once we recognize the two faces of intention, how should we think of the phenomenon of acting *with* an intention? This is the question I will turn to in the next chapter.