

KIRK LUDWIG

IMPOSSIBLE DOINGS*

(Received in revised form 16 July, 1991)

1. INTRODUCTION

This paper attacks an old dogma in the philosophy of action: the idea that in order to intend to do something one must believe that there is at least some chance that one will succeed at what one intends.¹ I think that this is a mistake, and that recognizing this will force us to rethink standard accounts of what it is to intend to do something and to do it intentionally.

We can begin by asking whether intention can be reduced to a belief that one will do what one intends, or a belief that one will do what one intends plus some desire to do it, then turn to the question whether some sort of belief that one will or can do what one intends is at least a necessary condition on intending to do something. I will argue by means of examples that this is not so, explain how it is possible and draw out some of the consequences for the received picture of action explanation, respond to some objections, and suggest the kind of account of intention that I believe will allow for the possibility of intending to do what one believes to be impossible.

2. IS INTENDING TO *A* A BELIEF THAT ONE WILL *A* OR A BELIEF THAT ONE WILL *A* PLUS SOME DESIRE TO *A*?

A natural first suggestion about the relation between intending to *A* and believing that one will is that intending to *A* simply *is* believing that one will *A*. It is easy to describe a case, however, in which one believes that one will *A* without intending to do so. In adhering to a high protein diet I promote the growth of the hair on my head. But I adhere to the diet for the sake of my nails, of which I am vain, and merely foresee the effect on my hair, for which I care not. Thus I believe that I will

promote the growth of my hair, without intending to do so. I may, again, believe that I will lose some money in Las Vegas without intending to do so. For I may intend to go to Las Vegas, and intend to gamble, and believe that if I gamble I will lose some money, although I do not intend to gamble *in order* to lose some money. Or I may believe that because many of my friends are poor judges of good wine it is inevitable that in the coming year, under the influence of conviviality, I will imbibe some sour concoction, although I neither want nor intend to do so. Not every belief one has that one will do something, then, is accompanied by an intention to do it. So an intention to *A* is not a belief that one will.

One thing missing in these cases is any desire to do what we expect to do. It is either unwanted, or something to which we are indifferent. It might be suggested, then, that an intention to *A* is not simply a belief that one will *A*, but, rather, a belief that one will *A* plus a desire to do so. However, even a desire to do what one believes one will do is not sufficient to intend to do it. For suppose that I know that I am rather absent minded, and altogether a creature of habit. I am in the habit of going to the barber every week. Financial straits force me to cut back. I do not want to, because I abhor long hair, but my barber is expensive, and paying the rent is more important. Since I know that I am absent minded, and that I allow my life to be governed by routine, I believe that come next Tuesday I will find myself in the my barber's chair as usual. But although I believe I will go to my barber's next Tuesday, and have a desire to do so, I do not now intend to have my hair cut next Tuesday.² So wanting to *A* and believing that one will *A* are not sufficient for one to intend to *A*.

A natural amendment to this last suggestion is to require not only that one desire to do what one believes one will do, but desire that more than anything else. If one desires to do something more than anything else, and believes that one will, one believes either that one will do it (a) unintentionally or (b) intentionally. (a) Consider the case in which I believe I will do something unintentionally and want to do it more than anything else. Suppose I want to find a million dollars more than anything else, believe that I will (because I have been told so by my astrologer), but believe that I will find it accidentally and only accidentally (also because my astrologer told me this). I don't believe I need to

undertake, plan, or intend any action to fulfill my desire to find a million dollars, though I believe I will and want nothing better. I do not then intend to find a million dollars: I simply wait for it to happen. Thus wanting to do *A* more than anything else and believing one will do it unintentionally is not sufficient to intend to do it. Indeed, believing one will do it accidentally is believing that one will not do it intentionally, and it would be surprising if this gave one an intention to do it. (b) Now consider the case in which I believe I will do something intentionally and want to do it more than anything else. Suppose that I desire more than anything else to become a millionaire through playing the stock market, but that at this moment I have absolutely no idea how to go about doing that, and can conceive of no possible action on my part that would help me toward that goal. I believe that I will play the stock market and thereby become through my cunning and sagacity a millionaire, only because I believe that I am among God's elect, and that God will reward me by giving me the ability to make a million dollars on the stock market. At the moment, there is nothing I need do, no action to which I need contemplate committing myself, no plan, however vague, I need formulate. It seems to me that I cannot be said in that case to intend to make a million dollars on the stock market, though I believe I will make a million dollars on the stock market, and will do so intentionally, and want this more than anything else. I believe that I will intend to do this *in the future*, but at the moment I am just waiting for things to happen. I am planning on it, but not planning it. So even believing that one will *A* intentionally, and wanting to *A* more than anything else, is not sufficient to intend to *A*.³

There is a further difficulty that each of the suggestions entertained so far share. They all presuppose an understanding of the concept of acting intentionally, and if we understand what it is to act intentionally only if we understand what it is to have an intention, none of these accounts could serve as a reduction of the notion of intention. The concept of an action is not the concept of any event involving one's body, for an event is not an action unless it is intended under some aspect. A successful reduction of the concept of intention could not then appeal to a propositional attitude that includes in its content essentially reference to an agent's future *actions*.

This suggests that intending to *A* is not a composite of belief and

desire, but a psychologically distinct state. Intuitively, an intention is a plan to do something, where this embodies a commitment to carrying out the plan, and this is something one can lack while believing that one will do something, and believing that one will do it and wanting to do it, and even believing that one will do it and wanting to do it more than anything else.⁴

3. CAN ONE INTEND TO *A* WITHOUT BELIEVING THAT ONE WILL *A*?

But if believing that one will *A* is not sufficient for intending to *A*, even in conjunction with a desire to *A* more than anything else, it may still be necessary for intending to *A*. Paul Grice argued for this in “Intention and Uncertainty,”⁵ on the grounds that it is misleading to say that you intend to do something when you are in doubt about whether you will succeed. To take an example, suppose I think it very likely that I will be struck by a bolt of lightning for impiety before the football match tomorrow.

A: I intend to go the football match tomorrow.

B: Sounds fun.

A: I doubt I'll be there.

B: I thought you said you intended to go?

A: Yes, but I think I'll be struck by lightning for impiety before tomorrow.

B: Then you ought not to have said you intended to go, but instead that you intended to go if you were alive tomorrow, or, even better, simply: I'd like to go, and hope to be able to, but am not sure I will be able to.

However, as Grice himself has brought out in so many contexts, what it is misleading to *say* is not always *false*. Consider some examples Davidson gives in “Intending”:

suppose a man is writing his will with the intention of providing for the welfare of his children. He may be in doubt about his success and remain so to his death; yet in writing his will he may in fact be providing for the welfare of his children, and if so he is certainly doing it intentionally.⁶

in writing heavily on this page I may be intending to produce ten legible carbon copies.

I do not know, or believe with any confidence, that I am succeeding. But if I am producing ten legible carbon copies, I am certainly doing it intentionally. (EAE 92)

These examples show that one can do something intentionally while being uncertain that one is doing it. If I can be doing something intentionally while in doubt about my success, as Davidson remarks, I can surely intend to do something while in doubt about whether I will do what I intend.⁷ For example, surely I can intend to climb Mount Everest, without being sure (or even confident) that I will succeed, as would be shown by my being unwilling to make any serious wager on my success.

My fault in saying I intend to go to the football match tomorrow, when I believe that I will most likely be struck by lightning before then, is that the possibility I have in mind that may prevent me from going is something out of the ordinary, and not something you would know about. So when I say that I intend to go the football match tomorrow, without adding anything about my unusual circumstances, I know that you will suppose that the circumstances of my intending are similar to those of most people who intend to go to the football match tomorrow, and that you will infer that I will most likely be there. You may even plan on it. Since you know that I know this, and conversations are cooperative enterprises, you will suppose that if my circumstances are not normal, I will indicate this to you. Hence your umbrage when you find out that I did not. I should have indicated in some way that my circumstances were not the normal ones so as not to mislead you about the likelihood of my being at the football match tomorrow. But this no more shows what I said was false than does the misleadingness of my saying, in response your asking whether John is sitting as usual in the parlor this morning, that he is not sitting in the parlor, when I see him standing in the parlor looking out the window.

This leaves it open that there is a weaker belief condition on intending to *A* than believing that one will. We can distinguish a number of different possibilities. (1) You do not believe that you will not *A*. (2) You believe that you can *A*. (3) You believe that your intention to *A* confers on your *A*ing a non-zero probability. (4) You believe that your intention to *A* increases the probability that you will *A*. (5) You believe that there is a non-zero probability that you will *A*. No doubt there are

other variations as well. Rather than consider each of these variations separately, I will argue that even the weakest of these requirements, (5), is not necessary in order to intend to *A*. Since each of the others entails (5), this will suffice to refute them all.

4. INTENDING THE IMPOSSIBLE.

In this section I will give some examples to make plausible to you that one can do something intentionally even though one believes it is impossible for one to do it. I reach the conclusion that one can intend to do something which one believes to be impossible by the premise (which I will defend against recent attacks in section 6) that if someone *As* intentionally, than he intends to *A*.

(1) First example. I have foolishly parked my car in the driveway of my neighbor, Mr. Jones, an unemployed steel worker. Unfortunately, I remember trying to start my car last night and discovering that the battery was dead. Mr. Jones, not especially pleased, knocks on my door at 8 a.m. asking me to move my car out of his driveway so that he can go pick up his unemployment check. "Well, Mr. Jones," I say, "I'm not really sure I can get the car started," and I start to explain about the battery. But in the meantime Mr. Jones has got me round the neck and has lifted me about a foot off the ground and is growling in his inimitable way, "No excuses, pipsqueak: Just move the car." As I recall Mr. Jones's history of mental illness, I do not argue the point. I do not, however, think that I have any chance of starting the car — zero probability — even though I really would like to get the car started. For I know that you cannot start a car with a dead battery. But I believe that if Mr. Jones can see that I'm sincerely trying to start the car, he will see that it won't start despite my efforts, and I will at least escape with my life. So I get in the car, put the key in the ignition, pump the gas a couple of times, smile obsequiously at Mr. Jones, who is standing by the car window, and turn the key. To my astonishment, the car starts, and I realize suddenly that it was the car I parked in my other neighbor's driveway that had the dead battery.

Did I start the car *unintentionally*? I don't think we would say that I did. But then there's only one option left: I started the car intentionally, despite believing that it was not possible for me to do so.

(2) Second example.⁸ We've been shooting baskets. You insist that I can make a basket from half court. I believe that you have an exaggerated opinion of my athletic prowess. Nonsense, I tell you. But you pester me. Finally, I'm fed up, and I say, "You're crazy if you think I can do it, and I'm going to prove it to you. I'm going to get out there and try as hard as I can to make a basket from half court, and you'll see that I just can't do it." I position myself at half court, and try as hard as I can to make a basket. It turns out that your opinion of my athletic prowess wasn't so bad after all, and I sink the basket — and not just accidentally. I make it because I was trying to. My trying to was a condition on showing you I couldn't do it *despite trying*. And that I am successful is not a fluke, not due to a sudden gust of wind, or a muscle spasm. I had underestimated my strength and hand-eye coordination.

Did I make the basket *unintentionally*? I think obviously not. Therefore I sank it intentionally, even though I didn't believe that I was able to.

(3) Third example. David and Goliath. Imagine that David as he goes out to face the giant Goliath believes that although it is important to defeat Goliath, there's absolutely no chance of beating Goliath in battle. Still, honor requires David to face Goliath and to try to slay him. So he gets his slingshot out and whirls it over his head and lets fly with his best shot, and, to his surprise, slays Goliath. Did David kill Goliath *accidentally* or *unintentionally*? No. So David killed Goliath intentionally despite believing that it was impossible for him to do it.

(4) Fourth example. In the examples above, the agent performed an action that was intentional under a description licensed by something that some movement of the agent's body caused to happen. But the same sort of case can arise with respect to movements of an agent's body.

Abel has been in an accident and lost the use of his right arm for some time, and he has become convinced that he will never regain its use. Unknown to Abel, on one of his routine visits to his doctor, an operation is performed while he is under anesthesia to restore to him the use of his arm. Subsequently, his doctor asks Abel to draw a picture of a duck on a piece of paper with his right hand. Abel insists to his doctor that he cannot do it, since he has lost the use of his right arm. His doctor insists though that he try. To please her, Able tries, closing

his eyes to concentrate, and succeeds in drawing a picture of a duck with his right hand, though he believes while he is doing it that he can't and is not doing it.

Surely it cannot be intelligibly maintained that his drawing of the duck was something he did *unintentionally*? If this is right, then Abel draws the duck intentionally, unaware that he is doing so, believing it to be impossible.

If it's correct to say in any of these four cases that the agent intentionally did what he thought it was impossible for him to do, then, provided that someone performs an action intentionally under a certain description only if he intends to do something under that description, we can conclude that the agent had an intention to do something which the agent believed that it was impossible for him to do.

We may note that a feature of these cases is that in each the agent was genuinely *trying* to do that which he believed was impossible. This is an important point to which we will return.

In the next section, I consider two arguments designed to show the conclusion just reached could not be correct. Seeing where these arguments go wrong will help us to see not just that the conclusion is correct, but how it can be. We will turn to the question whether one can *A* intentionally under description *D* only if one intended to *A* under the description *D* in the following section.

5. HOW IS IT POSSIBLE TO INTEND TO DO SOMETHING WHEN ONE BELIEVES ONE HAS NO CHANCE OF SUCCESS?

In this section, I want to consider briefly two arguments against the possibility of intending to do what one believes one cannot do which are suggested in Donald Davidson's work.⁹ If these arguments were successful, then we would either have to reject the intuitions in my examples or we would be forced to deny the assumption that one does *A* intentionally only if one intends to *A*. I think seeing where these arguments go wrong will help us to see how it is possible to intend to do what one believes to be impossible.

The first argument rests on the assumption that "the reasons an agent has for intending to do something are basically of the same sort as the

reasons an agent has for acting intentionally.”¹⁰ The reason an agent has for acting intentionally consists of a belief/pro attitude pair that rationalizes the action. If Davidson is right, then the reasons an agent has for forming an intention also consists of a belief/pro attitude pair, the same belief/pro attitude pair that rationalizes the action intended. One’s *A*ing intentionally can be rationalized by a belief/desire pair only if one can construct a practical syllogism that shows that there was something to be said for the action, that is, only if an argument of this sort can be constructed:

1. All actions of type *A* are good insofar as they are of type *A*.
2. This action has a chance of being an action of type *A*.
3. This action is good insofar as and to the extent that it has a chance of being of type *A*.

Now, if one believes that no action one can perform has any chance of being an action of type *A*, then one won’t be able to construct any practical syllogism of this form in support of it. Hence, one will not have reasons to *A* intentionally which would rationalize one’s *A*ing, and if one has reasons to intend to *A* only when one has reasons which would rationalize one’s *A*ing intentionally, one will not have reasons to intend to *A*. And if one cannot intentionally *A* without having reasons which rationalize one’s *A*ing, then it is plausible to suppose that one cannot intend to *A* without having reasons which would rationalize one’s *A*ing.

Once we have laid the argument out, we can see that it assumes precisely what the examples I gave above undermine: that one can *A* intentionally only if one believes that it is possible for one to *A*. Consequently, it cannot be used to show that our intuitions in those cases are suspect, nor undermine the step to the conclusion that one can intend what one believes to be impossible.

This shows that if what I have said about the examples above is correct, the relation between standard action explanations, and doing an action intentionally under a certain description, is more complicated than is usually assumed. One can do something intentionally, even though a standard action explanation of what one does is *not* available, even though, in Davidson’s terms, there is no *rationalization* of the

action. (I will suggest below, however, that rationalization is still in a certain sense fundamental to action explanation.)

The second argument focuses on the utility of doing what one believes one cannot do. Suppose that I am considering hitting a home run by means of swinging a bat. For the moment we can suppose that the only reason I am considering swinging the bat is as a means to hitting a home run. If I decide that by swinging the bat there is no chance of hitting a home run, then, even if I value hitting a home run, I will not value swinging the bat on that account, since I think I have no chance of hitting a home run by doing so. In terms of decision theory, I value hitting a home run, but think the probability of my hitting a home run given that I swing the bat is zero, and have no other reason to value swinging the bat. Therefore, the value I assign to swinging the bat is zero also:

$$V(S) = P(H/S) \times V(H) = 0 \times V(H) = 0$$

According to Davidson, an intention is an 'all out' positive evaluation of a way of acting. Where a way of acting can have no value because it has no chance of bringing about any desired end, there is no possibility of forming an 'all out' positive evaluation of it, hence no possibility of forming the intention to do it.

But while this argument might show that one cannot form an 'all out' positive evaluation of an action type valued only because it is thought of as a means to an end that in the circumstances it cannot bring about, it would not show that one could not form an 'all out' positive evaluation of action types one thinks one cannot perform, provided either one values them intrinsically, or that if one performed them, they would contribute to bringing about some state of affairs one values intrinsically. For example, it doesn't follow from the fact that the value I assign to, say, swinging a bat, when my only aim is to hit a home run, is the same value I assign to hitting a home run. Thus, as we can see above, although the value of swinging the bat, $V(S)$, may be zero because the probability that I hit a home run given that I swing the bat, $P(H/S)$, is zero, this is compatible with the value of hitting a home run, $V(H)$, being quite high, hence, with my forming an 'all out' positive evaluation of it. The argument shows, then, not that one cannot intend to do something that one believes it is impossible for one to do, but only that

no action one believes it is possible for one to do will derive any value from being a means to an end one believes it is impossible for one to reach.

It may seem that this response misses the point. For it might be objected that one *could not* intend to hit a home run *without* intending to do something which is a means to that end. I think this is right. It makes no sense to suggest someone intends to *A* if that person has adopted no plan which could be conceived of as designed to bring about *A*. (Indeed, one of the examples in section 2 relied on this point.) But this requirement is compatible with one's having adopted a plan to do something one thinks can't succeed.

We can see how this is possible by noticing that performing an action designed to bring about an end can have a value for an agent for reasons other than its contribution to bringing about that end. It is this fact that the examples in section 4 exploit. When I was trying, successfully, as it turned out, to start my car, I was undertaking to start the car not because I thought I would be successful, but because I wanted to please my neighbor, and thought I could do that by sincerely trying to start the car, hence sincerely embarking on a plan of action designed to get the car to start (if anything would). When I was trying, successfully, as it turned out, to sink a basket from 50 yards, I was not undertaking to sink a basket because I thought I had any chance of success, but because I wanted to show you that I could not do it, and to do this I had to sincerely try to sink the basket, hence sincerely embark on a plan of action designed to sink the basket (if anything would). Similarly, when David slew Goliath, he did not undertake that with any hope of success, but because honor required it of him. When Abel undertook to draw a duck with his (as he thought) useless right arm, he did so not because he thought he could succeed, but to please his doctor. So, in each of these cases, the agent undertook to perform an action of a certain type for reasons other than its possible success. But in each case the agent undertook an action or plan specially related to the action the agent thought he could not perform.

The key here is that the action undertaken is conceived of as *designed to bring about a certain end*, although in the circumstances the agent believes that it cannot succeed. An agent cannot intend to start a car by committing himself to an arbitrarily chosen action. I cannot

stand in my driveway and flap my arms up and down with the intention of starting my car. It is crucial that what I commit myself to doing is conceived of by me as designed to bring about a certain end. But this, the cognitive element in intention, is weaker than belief. It perhaps implies that I believe that there are circumstances in which the probability that the action I am committed to will bring about its end is non-zero. But it does not imply that I believe that in the actual circumstances the probability that I will bring about that end is non-zero. I want to suggest that the conditions for undertaking an action designed to bring about an end *A* are just those in which it is correct to say that one is trying to bring about that end. For example, in shooting from half court I am obviously trying to sink a basket. But in standing in my driveway flapping my arms up and down, I am not trying to start my car. If correct, this is support for the conclusion I will draw in section 6 that one intends to *A* if and only if one intends to try to *A*.

If intending to *A* and *A*ing intentionally do not require having a belief/desire pair that would rationalize my *A*ing, is it still the case that I must have some intention which is rationalized by a belief/desire pair? I have argued that one can intend to *A* without believing that it is possible for one to *A*, and hence that one can intend to *A* and, fortuitously, *A* intentionally, although there is available no rationalization for one's intending and *A*ing intentionally. If we grant, however, that one must still commit oneself to some action (or, more austere, to some undertaking) that is conceived of as a means to the end one believes one cannot attain, then there is still a rationalization for some action or undertaking of the agent, though not the one he believes himself unable to perform. For while we rescued from the decision theoretic argument above the conclusion that one can do intentionally and intend to do something one believes one can't, we did so only at the cost of allowing that the agent undertakes a plan designed to bring about that end. On pain of infinite regress, then, we must allow that there is some undertaking that the agent commits himself to which he believes is possible for him. (I describe what the agent commits himself to as an undertaking (which may be an action), to allow that in certain cases, such as the example involving Abel in section 4, what the agent commits himself to is not an action.) Thus, the possibility of rationalizing an action or undertaking under some description is still necessary for an

agent to *A* intentionally, although it is not necessary that every description under which an action is intentional have a corresponding rationalization.

6. OBJECTIONS: INTENDING, INTENDING TO TRY, AND THE SIMPLE VIEW.

An important objection to the argument I have given is to deny that the inference from *x* *A*ed intentionally to *x* intended to *A* is legitimate, on the grounds that in the examples I described, before the agent's success, it would be very odd for him to say that he intended to start the car, or sink the basket, or slay Goliath, or draw a duck. The agents in these cases would be much more likely to report their intentions by saying, "I intend to try to *A*, but I don't think I've got any chance of success." An agent may even say if asked prior to the action whether he intended to, e.g., make the basket, "No, but I intend to try."¹¹ So what we should say is that the agents did not in fact intend to do what they succeeded in doing, but only intended to try to do it. Thus, even if the examples show one can *A* intentionally while believing it to be impossible, they do not show one can intend to *A* while believing it to be impossible.

This objection counsels that we embrace a belief condition on intending at the expense of a tight connection between *A*ing intentionally and intending to *A*. The trouble, in the light of my examples, is that we cannot hold both.

The first point to make against this objection is the same one made against Grice earlier. It may be odd, or misleading, for an agent to describe his intentions by saying that he intends to *A* when he thinks he has not the slightest chance of success, but this does not by itself show that it is false. The oddness, and misleadingness of it, can be attributed to the fact that ordinarily people do not intend to do things that they think they have no chance of succeeding at, so that if someone announces an intention to do something, we assume that he, at least, thinks that it is possible for him to do it. Since he knows that we will assume this, and we know he knows it, if he doesn't explicitly cancel the implication, we feel entitled to assume that he means us to suppose that he believes he has some chance of success. So when it comes out that an agent who announced an intention to do something did not believe

he had any chance of success, we feel that he has been misleading us, and should have said instead that he intended to *try* to do it, for trying does not require success, and would not suggest the agent thought he would succeed. I suggest that the oddness, and misleadingness, of saying that one intends to do something one believes one can't do, though one undertakes to try, is all due to the standard conversational implicatures that accompany announcements of intentions, and is not due to its being false for one to say that one has an intention to do something even though one does not believe one has any chance of success. In support of this, we can note that we use this locution also when our subjective probability for our doing an act is much higher than zero and in fact higher than 50%. If I think I have a 60% chance of sinking a basket from the free throw line, and am pressed about whether I intend to make the basket, I may well respond by saying, "I intend to try to make it," and even, "No, but I intend to try to make it," in order to indicate that I am not certain of success. I may respond in this way even when my subjective probability is much higher, if in the circumstances announcing that I intended to sink the basket would suggest to my audience that I was even more confident than I was. There is, I think, no incompatibility in supposing that I also intend to make the basket in such cases. If I preface my announcement of my intention to try to sink the basket with "No," this is intended *not* to deny that I intend to make the basket, but rather to deny the standard conversational implicature of *saying* that I intend to make the basket, namely, that I think I will or that it is quite likely that I will. (Indeed, I will suggest something stronger below, that far from these being incompatible, each entails the other.)

This first point only neutralizes the linguistic evidence against the claim that in these cases the agents intended to do what they thought they could not do, by explaining it in a way that makes it compatible with the agents having intentions that they do not announce. What reasons are there to think that if I *A* intentionally, I must have intended to *A*?

(1) First, the simplest explanation of what licenses an action as intentional under a certain description *D* is that the agent of the action intended to do an action of the type specified by the description *D*. According to this view, the point of saying an action is intentional

under a certain description is to say with what intention the action was performed. It's clear that intending to *A* and then *A*ing (in the right way) as a result of that intention is sufficient for the action to be an action of intentionally *A*ing. So it seems that what licenses the description of the action as intentional under a certain description is a fact about its causal history, namely, that it was caused (in the right way) by an intention to do an action of type *A*. Complications of this picture (dubbed the 'Simple View' by Michael Bratman) need to be justified. I suggest there is no good reason to complicate it.¹²

One argument to show that descriptions of actions under which they are intentional do not always have corresponding intentions is given by Michael Bratman in "Two Faces of Intention."¹³ Bratman's principal argument relies on an example involving an ambidextrous game player. (Another argument Bratman gives relies on holding that intending requires a belief that one will do what one intends, and so cannot be used against the present claim without begging the question.) The ambidextrous game player is playing a game in which the object is to guide a 'missile' into a certain target. The player is playing one such game with each hand. The games are linked so that it is impossible to hit the target simultaneously in each game. Bratman argues that if the game player hits target 1, he does so intentionally, in which case if doing something intentionally entails intending to do it, he must have intended to hit target 1. Given the symmetry of the case, that means the game player must also have intended to hit target 2. But since the game player knows it is impossible to hit both target 1 and target 2, if he had an intention to hit target 1 and an intention to hit target 2, he would have had intentions which were inconsistent in the light of his beliefs. Thus he would be criticizably irrational. He would violate what Bratman has called the principle of agglomerativity: if you can rationally intend to *A* and rationally intend to *B*, you can rationally intend to *A* and *B*. The game player cannot, given his beliefs, rationally intend to hit target 1 and target 2. It seems clear that the game player need not be irrational in adopting his strategy, hence need not be guilty of having inconsistent intentions. Bratman concludes that the Simple View must be false.

I agree with Bratman that *if* the ambidextrous game player's hitting target 1 is intentional under the description 'hitting target 1', then the

Simple View is untenable.¹⁴ However, I don't think that it is right to say that the ambidextrous game player's hitting target 1 is intentional under the description 'hitting target 1'. Or rather, I think that this under-describes what is intentional about his action. For the intention that the game player begins with is the intention to hit one or the other but not both of the targets, that is, the content of the intention is: I hit target 1 or I hit target 2 but not both. So when the game player hits target 1, he hits it intentionally because he aimed to hit one or the other of the targets. But the description under which it is intentional is: hitting one but not both of the targets. We can bring out the incompleteness of the description above by asking whether it could be spelled out by saying that the game player intentionally hit target 1 *rather than* target 2. That's clearly false. So when pressed we will say something like this: the game player hit target 1 intentionally insofar as hitting target 1 is hitting target 1 or target 2 but not both. It is obvious why we do not actually say anything this awkward when it is clear from the context what the agent's intentions are. One of the reasons we are misled here is that we do not want to say that the game player hit target 1 unintentionally, for that would suggest that it was done accidentally, not in satisfaction of any plan of the agent's, and that is not so. Hitting target 1 satisfied the agent's intention to hit target 1 or 2, and so hitting 1 was not accidental. But what we miss is that there is another option here in addition to simply saying he hit target 1 intentionally, period. The game player in hitting target 1 intentionally hit target 1 or target 2, and there was nothing accidental or unintentional about *that*. Compare this with intentionally flipping a coin to produce heads or tails. My flipping heads is not, we feel like saying, unintentional (it would be misleading to say so), but I did not flip *heads* intentionally, but only insofar as that is flipping heads or tails. Likewise, I did not hit target 1 intentionally, but only insofar as that was hitting target 1 or 2.¹⁵

Bratman anticipates this response¹⁶ and argues against it by contrasting the case of the ambidextrous game player with two others. In the first, an agent aims at a target knowing it is target 1 or target 2 but not knowing which it is. In this case Bratman agrees the agent only hits the target intentionally under the description 'hitting target 1 or target 2' but not under the description 'hitting target 1'. In the second case, two targets, 1 and 2, are close together, and an inexperienced gunman is trying to

hit one or the other, being incapable of aiming sufficiently well to pick out one from the other. In this case also Bratman agrees that the agent would hit target 1 intentionally only under the description 'one or the other of the targets'. Bratman argues that the case of the ambidextrous game player should be distinguished from these two because the ambidextrous game player is "trying to hit each of the two targets," though not both. "[E]ach of the two targets separately guides my attempt to hit it."¹⁷ Thus, the crucial difference between the cases according to Bratman is that the ambidextrous game player is trying to hit target 1 and trying to hit target 2, while in the other cases the agent is trying only to hit one or the other of the targets.

This response, however, seems to me to fall into the same difficulty that Bratman originally raised for the Simple View. A constraint on rational intending is that one be able consistently with one's beliefs about the world to carry out all of one's intentions. This is dictated by the role intentions play in guiding our actions and further deliberation about what to do. If our intentions were not consistent, then our actions and deliberations about what to do would not be coherent. So, if the ambidextrous game player can adopt his strategy rationally, as it seems he can, he must be able to do it without having inconsistent intentions. However, the very same considerations that show that having inconsistent intentions is irrational seem to show that it is irrational to try to do things that are incompatible. What I am currently trying to do guides my action, and constrains my deliberations about what to do simultaneously and in the future. If I am trying to finish a paper, I cannot simultaneously coherently try to play a game of tennis, because in my case, at least, I know these two things (finishing a paper and playing a game of tennis) cannot be simultaneously realized. So I am criticizably irrational not only for intending incompatible things, but for trying to do incompatible things, for what I am trying to do also guides and constrains deliberation and action. Thus, if the ambidextrous game player in Bratman's example were genuinely trying to hit target 1 and simultaneously trying to hit target 2, he would be criticizably irrational in the same way he would be if he intended to hit target 1 and intended to hit target 2, knowing they cannot both be hit. Since the ambidextrous game player can undertake his strategy coherently, he need not be trying to hit target 1 and trying to hit target 2. This removes the

difference that Bratman cites between the case of the ambidextrous game player and the supposedly contrasting cases. It also indirectly supports my claim that the game player hits target 1 intentionally under the description 'I hit target 1 or target 2'. For it is plausible to hold that one does not do *A* intentionally under a description *D* unless one was trying to *A* under the description *D*. To avoid attributing inconsistency to the game player, we must say he is trying to hit one or the other of the targets but not both. Since he was not trying to hit target 1, he did not hit target 1 intentionally under the description 'I hit target 1'. It is compatible with this that what guides the game player's effort to hit one or the other of the targets is two separate things, the movement of target 1, and the movement of target 2, for this is a fact not about the agent's aims, but the mechanisms for carrying them out. I no more need to be trying to do two separate things in this case than when clapping my hands.^{18, 19}

To drive home the point, consider the following alteration of the example. I'm an ambidextrous bowler. One evening I'm bowling with the devil. Now the devil is a mean bowler, and he has the ability to destroy one's bowling ball by throwing his pitch fork at it. The strategy I adopt to foil the devil is to bowl with two balls at once, knowing that he has only one pitch fork. So I bowl with both hands intending, certainly, to get a strike, but not knowing with which ball I will succeed, because I don't know which ball the devil is going to zap. I also know that if the devil doesn't zap one of the balls, they'll collide and I'll gain nothing from it. Fortunately, the devil doesn't know this. Now do I intend to get a strike with ball *A* and intend to get a strike with ball *B*? Certainly not. I simply intend to get a strike with one or the other. But when ball *B* is zapped and I get a strike with ball *A*, did I intentionally get a strike with ball *A*? I think not. I intentionally bowled with both hands in order to get a strike with one or the other ball, but I did not intentionally get a strike with ball *A*. The same thing should be said about Bratman's example. To see this, notice that we can transform this example by redescription into a case very much like Bratman's example. Suppose with each hand I aim at a different point in the triangle of pins. Call one point target 1 and the other target 2. If I do not get a strike intentionally with ball *A* hitting target 1, then certainly I do not intentionally hit target 1 either. For I could not intentionally hit target 1 without intentionally hitting the pins with *ball A*, which I did not.

What has gone wrong in Bratman's case is that we imagine ourselves, as it were, split in two for the duration of the game. One of us, we say to ourselves, is trying to hit target 1, while the other is trying to hit target 2. But that's a mistake. There's only one of us and what we're doing is simply trying to hit one of the targets, just as when bowling with the devil I am trying simply to get one or the other of the balls down the alley for a strike, not both.

(2) A second positive argument to show that in the cases I have described the agents intended to perform the actions they successfully performed despite believing it impossible for them to do so relies on the oddness of saying "I am trying to *A*, but I don't intend to *A*." The simplest way to explain this is to say that "*x* is trying to *A*" entails "*x* intends to *A*." Then saying "I am trying to *A*, but I don't intend to *A*" would be self-contradictory. It might be objected that we can clearly try to do things we don't want to succeed at. But in some of these cases I think we want to fail to do what we intend to do, and others I think can be described as cases in which one merely goes through the motions of trying. An example of the first sort is a fighter who has been told his legs will be broken if he doesn't throw a fight, but intends to win his bout for his professional integrity, though he also wants to fail to carry out his intention because of the threat. He is trying to win, but does not want to succeed. He hopes that his opponent will prove stronger. But if his opponent is not stronger, he will not allow him to win. He intends to win the bout if he can, despite wanting not to (and hoping he will not) succeed. For an example of the second sort, we can imagine an executioner who has been charged with garroting his brother. He believes that his brother has an iron neck, and that consequently he cannot be strangled. He will lose his job, however, if he does not make his employer believe that he is trying to strangle his brother. So, as we say loosely, he tries to strangle his brother, without, of course, intending to do so. However, in this case, I do not think we should say that he is really trying to strangle his brother, but rather that he is merely going through the motions of doing so. For clearly, if it appeared to him that he was succeeding, he would immediately stop, come what may for his future employment. So he was not trying to strangle his brother, but to convince his employer he was. If you are genuinely trying to do something, success is not one of the obstacles to be overcome.²⁰

If I am correct that "I am trying to *A*" entails "I intend to *A*," then it

follows immediately from my examples that one can intend to do something one believes to be impossible, because in those examples it is clear that the agents are trying to do something that they believe to be impossible. But we can reach a somewhat stronger conclusion. If “I am trying to *A*” entails “I intend to *A*,” then “I intend to try to *A*” entails “I intend to *A*.” For if one intends to try to *A*, and carries out that intention, then one will have tried to *A*. From this it follows that one intended to *A*. But since there is no psychological difference between an individual who carries out the intention to try to *A* and one who fails to do so, if an individual who carries out his intention to try to *A* intends to *A*, then even if he does not carry out his intention to try to *A*, he intends to *A*. It seems to me clear that if someone intends to *A*, he intends to try to *A*. This gives us the biconditional: you intend to try to *A* if and only if you intend to *A*.²¹

It is no objection to this that one will be surprised if one succeeds in doing something one had thought one couldn't do. For one can plan something and still be surprised by one's success. This could be the case in Davidson's example above of someone trying to write ten carbon copies. He may think his chances of success quite low, though not zero, and be pleasantly surprised that he has succeeded.

Thus, in my view, the *Simple* is to be preferred to the *Baroque*.

7. PLANNING, WANTING, AND COMMITMENT.

Before concluding, I want to make a few brief remarks about what it is to have adopted or acquired a plan to *A*, or to plan to *A*, as opposed to simply having a plan in mind or entertaining a plan. I have suggested that to have an intention to *A* is to have formed a commitment to carry out a plan to *A*. If one intends to *A*, then one plans to *A*. But one does not, I think, plan to *A* unless one wants to *A* and is committed to *A*ing. It is this *commitment* to a plan of action, I believe, which constitutes the special role of intention as a propositional attitude. Thus, in my earlier examples, it is important that the agent in each case genuinely wanted to bring about the action that he believed to be impossible, and was committed to doing so. To take the third case as an example, if David and Goliath were secretly friends, but David felt bound by honor at least to go through the motions of trying to kill Goliath, still believing

it to be impossible for him to do so, then I do not think we would say that he killed Goliath intentionally, even though he killed Goliath, and was going through the motions of trying to do so. That this was not part of his plan or what he had committed himself to is shown by the fact that it was something that he did not want to happen (= it is not the case that he wanted it to happen) and had not committed himself to. It would be entirely appropriate for David to say that he did not intend to kill Goliath at all in this case, and did not kill him intentionally, in contrast to the case in which he did want and try to kill Goliath, but believed it was impossible to do so. The latter case is one of unexpected success; the former is an unexpected miscarriage of his plan, which was to save honor (or its appearance) by going through the motions of battling Goliath.

8. SUMMARY AND CONCLUSION.

I have argued that there is no essential connection between intending to *A* and even believing that it is possible for one to *A*. Intentions are clearly not identical with a belief of any kind directed toward the possibility that one will do some action, nor are they identical to any belief plus some desire, even if that desire is a desire to do some action more than any other. The reason it appears that there is some connection between intending to *A* and having a belief that it is at least possible for one to *A* is that generally the only reasons for adopting plans for action depend upon there being some likelihood that the plans will be successful. That is the point of the plan. However, this is at best a pragmatic connection between having an intention, or being committed to a plan for action, and believing that there is some possibility that one will carry it out successfully. It is possible to adopt a plan of action, to intend to do something, even though one thinks one has no chance of success, in circumstances in which there are other virtues in adopting or undertaking a plan for action. Thus it is possible that in undertaking to do something one believes is impossible one will succeed in doing it intentionally.

If this is right, then we must give up the view (prominent in Davidson's philosophy of action) that holds that there is an intimate connection between intending to do and doing something intentionally under a

certain description, and being able to explain what was done by citing a belief/desire pair that rationalizes the action under that description. Intending and doing intentionally depend on commitments which can develop independently of what Davidson has called primary reasons for what is intended or done.²² Rationalizing an action under some description *D* remains a necessary condition on its being an action; but not every description under which *A* is intentional and intended needs to have a corresponding rationalization.

If we accept that in the kinds of cases I described above an agent can intentionally do something which he believed he could not do, then my explanation of this possibility in terms of identifying intentions with commitments to plans for acting constitutes an indirect argument in favor of the view that intentions are commitments to plans for action. In any case, if we accept the intuitions in these examples, a constraint on any acceptable theory of intention is that it explain how it is possible for an agent to have an intention to do something when he believes that it is impossible for him to do it.

NOTES

* I would like to thank the following people for much helpful discussion of the issues this paper takes up: David Chan, Tony Dardis, Michael Della Rocca, Ariela Lazar, Rob Marsh, Alfred Mele, Leo Van Munching, Dugald Owen, Greg Ray, Kwong-Loi Shun, Tom Smith, and Bruce Vermazen. I would also like to thank an anonymous referee of this journal for helpful comments on the penultimate version of this paper.

¹ This thesis is widely defended, often with a stronger belief requirement than this, but not universally. Hugh J. McCann, in an article I read after completing this paper, argues against the view that we must at least not believe we will not do what we intend to do. On a natural construal this is to argue against the view that we must at least regard it as more likely than not that we will do what we intend to do. At one point though he also suggests one might believe one's chances of success zero but still intend to do something. See McCann's "Rationality and the Range of Intention," in *Midwest Studies in Philosophy*, X, edited by Peter A. French, Theodore E. Uehling, and Howard K. Wettstein (Minneapolis: University of Minnesota Press, 1986), pp. 191–211.

² I believe this provides also a counterexample to Wayne Davis's somewhat more complicated analysis of intention: *S* intends that *p* iff *S* believes that *p* because he desires that *p* and believes his desire will motivate him to act in such a way that *p*. See his "A Causal Theory of Intending," *American Philosophical Quarterly*, 21: 43–54.

³ This point should extend to 'all out' positive evaluations in favor of an action of a certain type. If this is right, Davidson's view that intentions are 'all out' positive evaluations of a certain type of action cannot be correct.

⁴ Thus I am in basic agreement with Michael Bratman about the best way to think about intentions, namely, as plans or partial plans, although, as we will see, there are other aspects of Bratman's theory of intention I am in disagreement with. See his

Intention, Plans, and Practical Reason, (Cambridge, Massachusetts: Harvard University Press, 1987). I do not claim so far to have established that intentions are plans. My claim in the end will be that thinking of intentions as plans is one way of understanding how it is possible to intend to do something while believing it is impossible for one to do so. I regard the possibility of intending to do something without believing it possible to provide indirect support for Bratman's account of intentions as plans or partial plans.

⁵ H. P. Grice, "Intention and Uncertainty," *British Academy Lecture*, (London: Oxford University Press, 1971). Grice's position is complicated by his arguing that there are reasons to regard the attitude that accompanies intention as something other than belief. But he adheres to the claim that intention involves something at least like acceptance that one will do what one intends. This strong belief requirement finds a qualified endorsement in Gilbert Harman's "Willing and Intending," in Richard Grandy and Richard Warner, eds., *Philosophical Grounds of Rationality* (Oxford, 1986), pp. 363–380.

⁶ Donald Davidson, "Intending," in *Essays on Actions and Events*, (Oxford: Clarendon Press, 1980), pp. 91–92. Cited below as 'EAE'.

⁷ Although Davidson here denies that one must believe that one will do what one intends, he does endorse a weaker belief condition on intention, namely, that one believe that one can do what one intends. See "Intending," in EAE pp. 91–2, and "Reply to Pears," in Bruce Vermazen and Merrill B. Hintikka, eds., *Essays on Davidson: Actions & Events*, (Oxford: Clarendon Press, 1985), pp. 211–15 (cited below as 'ED').

⁸ Rob Marsh suggested this example to me.

⁹ "Reply to Pears," in ED, pp. 213–4. Both of the arguments I discuss I think are suggested in the passage referred to, though they are not distinguished as different arguments.

¹⁰ "Reply to Pears," in ED, p. 213.

¹¹ Al Mele gives examples of this sort in "Intention, Belief, and Intentional Action," (*American Philosophical Quarterly*, 26, (January 1986): 19–30) and "She Intends to Try," (*Philosophical Studies*, 55, (January 1989): 101–106). In the first of these papers Mele argues for a functional account of intention which does not see a belief requirement as at the core of intention. The belief requirement that one not believe that one will not do what one intends is motivated solely by appeal to usage.

¹² Fred Adams has argued, persuasively, I think, that so far suggested replacements of the Simple View do not have the resources to adequately answer the two principal questions the Simple View is designed to answer: (1) When is an action intentional under a certain description *D*? Answer: when the agent of the action had an intention to do an action of type *D* which caused the action in 'the right way'. (2) When is an action not intentional under a certain description *D*? Answer: when the answer to question (1) does not apply. See his "Intention and Intentional Action: The Simple View," *Mind & Language*, 1 (Winter 1986): 281–301.

¹³ *The Philosophical Review*, 93 (July 1984): 375–406. The argument is repeated with replies to some objections in Bratman's *Intention, Plans, and Practical Reason*, chapter 8.

There are other arguments against the Simple View, but I regard this as the most powerful. One strategy in arguing against the Simple View is to hold that habitual actions are done intentionally, but in the absence of having formed an intention to do them. The response to this is to distinguish between *having* and *forming* an intention, where this is understood to mean arriving reflectively at an intention, and to insist that the argument conflates these. Another strategy is to argue that accomplished actions have parts which we do intentionally although we do not have intentions to do them. I agree that what is done as a part of an intentional action, such as moving the index

finger in a certain characteristic way when tying one's shoes, is done intentionally, but, I think, only under a less fine-grained description than is suggested by the detailed description of that part of the action. A third strategy would be to hold that sometimes we can be said to do something intentionally when we see it as an inevitable or constitutive part or consequence of what we do, even though we do not particularly want to do it or even want not to do it. Bratman gives an example of a runner wearing down the soles of his shoes. Bratman says he wears down the soles of his shoes intentionally, because he intends to run a race wearing them and realizes that this will involve necessarily wearing down the soles of his shoes. I do not find such cases compelling. There are differences between the runner who intentionally wears down the soles of his shoes and the runner who does not, when both believe it a necessary consequence of running the race wearing their shoes. The former, but not the latter, e.g., would be inclined to take action to avoid having the soles of his shoes treated to prevent wear. For my purposes in this paper, however, it is enough to note that none of these ways of licensing an action as intentional under a description *D* without there being a corresponding intention to do an action of that type applies to my examples.

¹⁴ One might object that Bratman's principle of agglomerativity is not internal to intending, but is instead merely a pragmatic constraint on intending. I will not take this line here. Nor will I try to offer a defense of the principle of agglomerativity against this objection. I will try to show that even if we accept the principle of agglomerativity, we do not have to give up the Simple View.

¹⁵ Here is another way of thinking about this. Suppose we reconceptualize the situation as one in which the game player is attempting to guide two different missiles into two different parts of one target. And suppose it is still the case that he knows (or at least believes) that it is impossible to strike the target with both missiles. He believes, nonetheless, that his chances of a hit on the target are best if he guides in one missile with each hand. Suppose he hits the target on the right side. It is clear that he hit the target intentionally, but did he hit the target on the right side intentionally? I think it is clear that he did not. But how is this significantly different from the example as described by Bratman? The bowling game with the devil described below in the text is intended to bring home the same point.

¹⁶ *Intention, Plans, and Practical Reason*, pp. 116–119.

¹⁷ *Intention, Plans, and Practical Reason*, p. 118.

¹⁸ Bratman recognizes that he has to treat tryings or endeavorings differently from intendings. However, it seems to me that his principal motivation for doing so is his conclusion in the case of the ambidextrous game player, and the way seeing the ambidextrous game player as trying to hit target 1 plays a role in his defense of his conclusion in that case. Since the difference between tryings and intendings is motivated by appeal to a case which is undermined if there is no difference, the case cannot be appealed to in support of the difference without begging the question.

Bratman does suggest that it is possible to have guiding desires that are not subject to the same constraints that intentions are. But I do not find the case persuasive. Suppose Elizabeth wants both to marry Robert and to marry Francis. Bratman argues that Elizabeth may form two subplans, one to marry Robert, guided by the desire to marry Robert, one to marry Francis, guided by the desire to marry Francis, with the idea that the world will decide which plan will be successful, and not be criticizably irrational. This does not seem to me to be so. Elizabeth may surely carry out negotiations designed to bring about marriage between her and Francis and her and Robert, but what guides her negotiations cannot be desires heading toward incompatible ends if she is not to be irrational. As long as desires don't enter the arena of action, their incompatibility is not criticizably irrational. Once they take on a role in guiding action directly, if that is possible, they must be subject to the same constraints that apply to intentions. For those constraints come simply from the need that the world one sees one's efforts as directed toward bringing about itself be consistent.

¹⁹ I would like to note here an interesting objection to my account due to Tom Smith. Consider the example of David and Goliath. David believes it impossible to kill Goliath with a slingshot, but thinks honor requires him to make the attempt. On my account David intends to kill Goliath. Let us modify the example slightly. David believes he will fail to kill Goliath with the slingshot, but wants to kill him, so he brings along a gun which he intends to use to kill Goliath after the slingshot fails. Thus it seems my account would have us attribute to David two intentions, the intention to kill Goliath with the slingshot, and the intention to kill Goliath, a bit later, with the gun. Not even Goliath is big enough to be killed twice. We may suppose David knows this. So it looks as if in taking my line we will be forced to attribute to David inconsistent intentions, and so run up against the principle of agglomerativity. The appropriate response, I think, is to see the second intention, the intention to kill Goliath with a gun, as a contingency plan. In a modified example, David might think it possible to kill Goliath with a slingshot, though unlikely, and so intend to kill Goliath with a slingshot (still a requirement of honor), but also intend, if he should fail, to kill Goliath with the gun. In other words, the second intention is a conditional intention. It kicks into action on condition that the first intention fails to be successfully carried out. The only difference between the first case and the second case is the degree of conviction David has that his contingency plan will have to be put into action.

²⁰ Additional support for this entailment is provided in McCann's essay, "Rationality and the Range of Intentions." McCann focuses on the question: what distinguishes trying to *A* from doing other tryings or actions which the motions that are trying to *A* might have been? The most straightforward answer is that the agent intends to *A*, and it is difficult to see what the alternative is suppose to be. It is not a trying at all unless there is an intention in the picture. If it is not the intention to *A*, what is it?

²¹ What trying is obviously deserves more discussion. But that is the work of another paper.

²² See Davidson's "Actions, Reasons, and Causes," reprinted in *Essays on Actions and Events*, (Oxford: Oxford University Press, 1980), 3–19.

Department of Philosophy
University of Florida
240 A Daur Hall
Gainesville, FL 32611–2006
USA