5
Intentional Action

5.1 Introduction

A causalist account of intentional action gives causation a necessary role to play. The caricature has it that an agent acts intentionally only if her behavior is caused in the right way by a relevant mental state (or that state's) acquisition. What more is required is a matter of much debate—there are many varieties of causalist account.[1] A shared central flaw is the lingering worry regarding deviant causation.

We should be feeling okay. We solved that problem.

In this chapter, I leverage accounts of control and non-deviance to offer a new account of intentional action.

I say new. But the account has much in common with, and is inspired by, earlier causalist accounts (e.g., Goldman 1970; Brand 1984). Some of this will become apparent in the next section. I begin to explicate my own account in conversation with the account offered by Mele and Moser (1994).

5.2 Intentional Action: Mele and Moser

Mele and Moser's account is still, twenty-five years on, the most detailed and to my mind convincing analysis of intentional action. The paper that sets out the analysis has been influential for my own thinking on intentional action. Further, some of the merits of my account are more easily seen via engagement with this account. For an account of control helps to undergird some of the points Mele and Moser make. Finally, the differences that emerge may prove instructive for those invested in the project of offering a causalist account (or explication) of intentional action. I do not doubt many

[1] There are also many discussions of the many varieties of causalist account. For good introductions, see O'Brien (2015), or essays in Aguilar and Buckareff (2010), or in D'Oro and Sandis (2013).

The Shape of Agency: Control, Action, Skill, Knowledge. Joshua Shepherd, Oxford University Press (2021).
 DOI: 10.1093/oso/9780198866411.003.0005

will disagree with some of the judgments my account embeds. The discussion will serve at least to make my judgments plain and to offer readers the chance to explicitly diverge at various choice points.

Here is the final analysis Mele and Moser (1994: 63) offer:

> Necessarily, an agent, S, intentionally performs an action, A, at a time, t, if and only if:
> (i) at t, S A-s and her A-ing is an action;
> (ii) at t, S suitably follows—hence, is suitably guided by—an intention-embedded plan, P, of hers in A-ing;
> (iii) at the time of S's actual involvement in A-ing at t, the process indicated with significantly preponderant probability by S's on balance evidence at t as being at least partly constitutive of her A-ing at t does not diverge significantly from the process that is in fact constitutive of her A-ing at t; or (b) S's A-ing at t manifests a suitably reliable skill of S's in A-ing in the way S A-s at t; and
> (iv) the route to A-ing that S follows in executing her action plan, P, at t is, under S's current circumstances, a suitably predictively reliable means of S's A-ing at t, and the predictive reliability of that means depends appropriately on S's having suitably reliable control over whether, given that she acts with A-ing as a goal, she succeeds in A-ing at t.

Mele and Moser arrive at this analysis via consideration of a wide range of puzzling and difficult cases. In order to illuminate the zones of agreement and disagreement, I will offer commentary on aspects of each of their four conditions, with discussion of some relevant cases.

First, this is an analysis of what it is to intentionally perform an action. As condition (i) makes clear, this leaves unanalyzed the notion of action. That is something I would rather avoid. I do not speak, then, of actions that are intentional. Instead I speak of behavior that amounts to intentional action. The notion of behavior is, in my view, exceedingly broad. To call some event a bit of an agent's behavior is not to say the agent has done anything in any significant sense. Anything within the scope of an agent's causal potency could be construed as her behavior. The interesting thing is that some bits of an agent's behavior qualify as intentional action.

Condition (ii) speaks of an agent's suitable following of a plan. Mele and Moser do not offer an account of following. But they draw a connection between following and guidance by a plan, and between guidance and deviant causation cases. This is plausibly the identical, or at least a similar, connection

to the one between control and non-deviance. It is thus not implausible to offer an account of following, and hence guidance, in terms of control:

Plan Following. An agent J suitably follows—hence, is suitably guided by—a plan P (or certain aspects of P) when J exercises control in bringing her behavior to approximate (to some sufficient degree) the representational content of P (or certain aspects of P).

Condition (iii) is disjunctive. The first part mentions an agent's evidence. The second part mentions an agent's suitably reliable skill. I will discuss the second part first. Mele and Moser do not offer a detailed discussion of skill, but do offer the following useful characterization:

> At time t, an agent, S, is skilled to some non-zero degree at A-ing in manner M if and only if S, at t, has a propensity to A in manner M, given that S has a corresponding intention in a situation suitably accommodating for S's A-ing in manner M, and that propensity exceeds any provided by mere chance given such an intention and situation. (Mele and Moser 1994: 57)

This characterization appears congenial to my account of control. In some respects, my account seems simply a more involved working out of various details—what it is for a situation to be suitably accommodating, for example, and how best to think of degrees of control beyond zero. But one might pause over their language of a propensity that exceeds any provided by mere chance. Mele and Moser comment: "To have a high propensity to roll non-doubles with a pair of fair dice (simply by rolling the dice) when one has a corresponding intention is not to be skilled at rolling non-doubles. Such a propensity is predicted by mere chance" (57).

This claim has the potential to create a problem for my account of control. To see why, consider the following case from Mele and Moser, who credit Michael Costa:

> Mike, a normal person, is playing a game with a pair of fair dice. He will win $20 on his next roll if and only if he throws something other than "boxcars" (two sixes). Mike, wanting to win, has a simple plan: he will throw a non-boxcar roll and win the money. Mike realizes that there is a slight chance that he will roll boxcars, but this does not threaten his plan. As it happens, he throws a seven. (Mele and Moser 1994: 62)

Given that the dice are fair, one can imagine that Mike has very close to a 92 percent probability of rolling non-boxcars. On my account then, doesn't Mike have a high degree of control with respect to intentions to roll

non-boxcars? And isn't this an odd thing to think? After offering the case, for example, Mele and Moser make this comment: "Mike lacks a kind of control over the dice required for his intentionally throwing a non-boxcar roll" (62). Intuitions may side with Mele and Moser here, suggesting that my account does violence to the notion of "control over."

Mele and Moser's decision here is to exclude cases like these—cases of what they call statistically reliable luck—from their account of intentional action. In so doing, they explicitly contrast control with statistically reliable luck. And in their condition (iv), they claim that an agent's following a "suitably predictively reliable means" of A-ing depends upon her "having suitably reliable control over" whether she succeeds.

One might wonder, however, what reliable control amounts to if it is not simply a high propensity to succeed. Regarding Mike and the dice, Mele and Moser comment that "Mike has no control over which sides land face up. He thus has no control over whether, given his throwing the dice, he throws a non-boxcar roll (as opposed to a boxcar roll)" (62). This suggests a requirement of something like a power of influence regarding certain aspects of one's behavior. How to think of this power of influence? A power to determine every aspect of one's behavior is far too strong.

Some have thought of this power in terms of positively shifting the probability of success (Ross 1978; Costa 1986). But Joshua Gert (2004) has argued convincingly against this kind of view. For one can construct cases of obvious intentional action on which agents lack the ability to shift the probability of success—just imagine an agent triggering a device that has a 92 percent probability of killing another person.

I doubt there is a good way to finesse the intuition behind dice roll cases.[2] But let's try. Consider Nash, who is very good at shooting free throws. Nash makes 92 percent of the ones he shoots. One might think that the 8 percent are in some sense Nash's fault. In those cases, Nash failed to exercise his abilities properly. He could have made the shot, but he didn't. Ask Nash about this, however, and he might disagree. Nash might admit that some subset of the misses are his fault. He might say this:

Nash's Speech. Look, sometimes I know why I missed. And sometimes I can intervene to change the probabilities in my favor. I can intervene on the

[2] Gert suggests that dice roll cases are not cases of intentional action because the agent's plan "does not include a representation of the fact that the dice come up [non-boxcars] because she rolls them" (Gert 2004: 155). But it seems clear an agent's plan could include such a representation—indeed, in Mele and Moser's case the agent throws the dice in order to roll non-boxcars in order to win money.

processes constitutive of my behavior. So, in some cases I think I might have done better if I had taken a deep breath, focused, tried harder, removed distractions, visualized my plan more clearly, or whatever. But even having done these things, sometimes I get to a place at which I have done everything I can do, and the probabilities are what they are, and they are short of 1. Sometimes I just miss. I'm good, but I'm not perfect. I think that my missed free throws, sometimes, are not the result of any failure on my part—at least no failure I can do anything about. I think sometimes my miss is the result of a normally very reliable process misfiring. Or, possibly, the process underlying my shots has margins that occasionally fall outside what is required for success.

I'm on Nash's side. I think reliability is the most we can require. If so, I see no obvious reason to exclude processes like those involved in the dice roll. After all, if Nash takes up a dice-throwing game that requires frequent non-boxcar rolls, he might plausibly come to view his situation as similar to the free throw situation. It takes no practice for Nash to be just as good at non-boxcar rolls as free throws, of course, but that is an inessential detail.

I admit, mind you, the unintuitive nature of this response in Mike's case. The rigidity of the probabilities, the opaqueness of the details to the agent that initiated them, and the agent's inability to intervene on these processes all play with our intuitions, tempting us to think of what is happening as unusual or somehow deeply non-agentive. But I submit that most intentional action includes aspects that have rigid probabilities, that have opaque details, and that do not permit intervention. This is something we should simply accept. If we do, then we can understand "suitably reliable skill" in condition (iii) simply as sufficient control, and we need not separate the skill of condition (iii) from the control of condition (iv).

The first part of condition (iii) mentions an agent's evidence. Why think an agent's evidence is relevant to whether she has intentionally acted? Mele and Moser offer this case:

> Ann works as an admissions supervisor at an orchestra hall. She gives red admission slips to women and blue slips to men. Since the orchestral concerts are formal occasions, Ann always wears white gloves while working. Without examining the slips, she efficiently gives the men blue slips from her left pocket, and the women red slips from her right pocket. All goes as planned until Ann hands Harry an admission slip from her left pocket. Unbeknown to all concerned, Ann's admission slip for Harry is actually a white piece of litmus paper that instantaneously turns blue when touched

> by a human hand. Harry's hand turns Ann's white litmus paper blue. Consequently, Ann gives Harry an appropriate admission slip.
>
> (Mele and Moser 1994: 52)

What did Ann do wrong? Mele and Moser diagnose the source of the coincidence as stemming from a divergence between the route Ann took and her evidence regarding that route. Now, I agree that an agent's evidence can be relevant to their behavior. But I suggest that this relevance runs through the notion of control. Insofar as agents rely on their evidence, they need good evidence. Misleading evidence will tend to lead to failure, rather than success. The way it does so is by infecting an agent's plan for action. Agents acting on bad evidence are agents acting on a bad plan—a plan that embeds bad expectations, that gives improper directions. Bad plans succeed less often than good plans. That is, bad plans lower an agent's degree of control. So I suggest we assimilate the condition regarding evidence into a broader condition regarding plan quality and control.

Condition (iv) requires further attention. Mele and Moser maintain that the route the agent takes in executing her plan must be a suitably predictively reliable means to success. They link this requirement of predictive reliability to the agent's control, where the control operative here is a notion I have challenged. But there may nonetheless be something in this notion of predictive reliability. They explicate it as follows:

> Intentional action, on the intended interpretation . . . requires that—given just (a) S's suitably reliable nonmisleading evidence [that is, evidence devoid of false propositions] concerning whether she will A at t in her present circumstances, and (b) knowledge concerning what sort of reliable skill, if any, S has with respect to A-ing at t in her present circumstances—a conceiver who understands all the relevant concepts (sufficiently to wield them in any prediction involving just those concepts) could reasonably predict that (the route followed in) S's attempted execution at t of an intention incorporating the relevant action plan P will result in her A-ing at t. (Mele and Moser 1994: 60)

I have challenged the necessity of an agent's evidence for her ability to act intentionally, except as it contributes to the quality of an agent's plan, and thereby her control. Beyond this, whether an outside observer who has relevant knowledge could use the agent's route to predict whether she will succeed can be understood as a way of honing in on the agent's own success-rate

in behaving as she does. That is, we can understand the requirement of predictive reliability in terms of the possession of control.

5.3 Intentional Action

As I said above, I already found Mele and Moser's proposed analysis of intentional action compelling. It is possible to undergird elements of that analysis with the account of control I have developed, deepening our understanding of the nature of notions like plan following, and affording a response to the most long-standing complaint against causalism about intentional action, namely, that deviant causation undermines any causalist account. Further, if one sides with me on certain choice points in the analytical tree, one may be disposed to prefer a more parsimonious account of intentional action.

Before I offer that account, additional commentary is required. As I noted, I am not offering an account of action that qualifies as intentional action. I cannot help myself to an unanalyzed notion of action, nor of the content of the action at hand—that is, of an A-ing. I am accounting for behavior that qualifies as intentional action. How would some bit of behavior come to do so?

By entering into relationships of sufficient approximation, or resemblance, with other key relata. The relata are the agent's behavior, (aspects of) the agent's relevant psychological state(s), and the action-type of which the action under inspection is a token. The idea is that behavior qualifies as intentional action only if the behavior sufficiently approximates an aspect of the relevant plan, as well as the action-type of which it qualifies as a token.

I am not, note, putting much weight on the notion of an action-type. This is a way of capturing a couple of facts. First, there are many ways to describe an agent's behavior as an intentional action. Sometimes some bit of behavior can be described as multiple different intentional actions. Sometimes it doesn't much matter—one or two or three different descriptions could be equally accurate. Second, it is arguably possible to describe some bit of behavior as an intentional action even if the agent was not conceptualizing the behavior as such (see the discussion below, in section 5.5.2). Here we impose an action-type on the behavior. If we are to do so, the imposition should be roughly accurate, and I assume we have a rough way of accurately deploying action descriptions.

Consider an uncontroversial case. LeBron intends to go to Los Angeles. He has a plan for doing so. It involves catching a flight from Cleveland. Suppose that LeBron's behavior is caused by the acquisition and persistence of this intention, and that LeBron exercises sufficient control over his behavior—he tells his driver which airport to go to, he finds the right plane. LeBron has intentionally done several things—told his driver where to go, caught a flight from Cleveland, gone to Los Angeles. Why? Consider:

It is important that LeBron's behavior conformed sufficiently to his plan, or to relevant aspects of it. If LeBron told the driver the wrong airport, and found the wrong plane, but strong gusts sent the plane off course and to Los Angeles, we might well think LeBron's arriving in Los Angeles unintentional.

It is important that LeBron's behavior conformed sufficiently to the action-type(s) of which the actions under consideration are tokens. The action-type going to Los Angeles is broad, but LeBron's behavior falls under it as a token. So, too, with LeBron's catching a flight. If LeBron announced an intention to go to Los Angeles, after which insidious agents gave him spiked milk, and transported him sleeping to the city of angels, we might have a question about whether this is a valid way to intentionally go to Los Angeles.

Finally, it is important that LeBron's plan was well-constructed in the sense that if LeBron followed it—if he exercised a sufficient degree of control in bringing behavior to approximate the plan—he would indeed bring behavior to approximate the relevant action-type.

So we arrive at the following proposal:

Necessarily, an agent, J, intentionally performs an action, A, at a time, t, if and only if:

(i) at t, J's behavior B is caused by a relevant plan-state P of J's (or its acquisition and persistence), or a package of such states, P*;

(ii) B sufficiently approximates a relevant aspect of the plan;

(iii) B approximates a relevant aspect of J's plan because J exercised control over B;

(iv) B sufficiently approximates the action-type of which A is a token;

(v) J has a good plan: J's following the plan, by exercising a certain degree of control in bringing B to sufficiently approximate the plan, is itself a reliable method, in the relevant comprehensive set of circumstances, for sufficiently approximating the action-type of which A is a token.

5.4 The Account Paraphrased

That account of intentional action took most of the book until now to develop. It contains terms the full explication of which date back to chapter 2, and include parts of later chapters. So it may be useful to state the ideas behind the account in plain language.

Intentional action is, in essence, the exercise of a sufficient degree of control in bringing behavior to approximate a good plan. It is composed of the following elements. First, a plan for action. Second, that the plan be good—that following the plan be a good way to satisfy goals embedded in the plan. Third, control over behaviors required by the plan. Fourth, a causal pathway that includes the psychological state(s) that represent the plan. Fifth, a relationship of approximation between controlled behavior and whatever action-type we use to classify the controlled behavior as intentional action.

5.5 Ancillary Issues

I may be able to further illuminate this account's commitments and (de) merits by discussing cases often pressed for or against one or another account of intentional action.

5.5.1 Side-effects

Gilbert Harman (1976) offers the case of a sniper who intends to kill a soldier, and who is aware—without intending—that in firing his rifle he will alert the enemy. He fires the rifle, kills the soldier, and alerts the enemy. Does he alert the enemy intentionally?

Some have said yes; many no. Mele and Moser argues as follows:

> Since the sniper does not unknowingly, inadvertently, or accidentally alert the enemy, it is natural to insist that he does not unintentionally alert the enemy. Such insistence does not entail, however, that the sniper intentionally alerts the enemy. There is a middle ground between A-ing intentionally and A-ing unintentionally. We locate "side-effect actions" of the kind in question on that ground. Insofar as such actions are not done unknowingly, inadvertently, or accidentally, they are not unintentional. Insofar as the agent is not aiming at the performance of these actions, either as ends

> or as means to (or constituents of) ends, they are not intentional either. We shall say that they are nonintentional. The ordinary co cept of intentional action requires the agent of an intentional A-ing to be aiming at A-ing. (Mele and Moser 1994: 45)

My account would not accord intentionality to the sniper's alerting the enemy. That was no part of the sniper's plan. I agree with Mele and Moser and many others about that. Whether a notion of nonintentional action is available is a separate matter, outside the parameters of my account. One might offer an account of such a notion, making it derivative on intentional action. I have no plan to do so. I would be just as happy saying the sniper knowingly caused the alerting of the enemy, without saying anything about action. That might give us all we could need (e.g., we can blame agents for knowingly causing things as well as for intentionally doing things). I might even add that an agent capable of knowingly causing A is an agent capable of intentionally A-ing. That does not render the notions equivalent, but it does indicate an important connection. If an agent knows that something is within her causal remit, then she can include that thing in her plan as a means or an end. And if she plans to bring it about, and succeeds in the way my account indicates, then it qualifies as an intentional action.

5.5.2 The Simple View and the Single Phenomenon View

One ongoing debate in action theory concerns a question about the relation between intentional action and intention (or, more broadly, the agent's plan-state) (see, e.g., Bratman 1984; Adams 1986; McCann 1987; Nadelhoffer 2006; McCann 2011; Wasserman 2011; Amaya 2018). Here are the two views:

Simple View (SV). An agent cannot intentionally A unless she intends to A.

Single Phenomenon View (SPV). In order to intentionally A an agent must execute some relevant intention, even if not an intention to A.

While SV presents initially as intuitive, it has been undermined by a range of cases, as well as by experimental philosophy suggesting that laypeople's philosophical judgments are better explained by SPV (Nadelhoffer 2006). For cases, see Bratman (1984), Di Nucci (2009), or this case from Mele:

> Alice is mowing her lawn—intentionally, of course. In the process, she has taken many steps. It would be strange—and, I believe, mistaken—to maintain that her taking her next routine step while pushing the mower is not an intentional action. But there is no need to suppose that Alice has an intention specifically to take that step. Given that she intends to mow her lawn at the time, is a proficient mower, encounters no obstacles requiring alteration of her gait, and so on, her mowing intention can do its work without her having a series of intentions corresponding to each routine step. (Mele 2005: 150)

I side with proponents of the SPV. The present point is that my account of intentional action permits latitude on the matter. I require sufficient approximation between plan-state and behavior. Depending on how one reads the SV vs. SPV debate, one might want sufficient approximation to amount to the Simple View or the Single Phenomenon View. If one feels the pull of Mele's above case, one might think behavior need not be exactly in line with the plan. One's plan is to mow the lawn, and that plan includes a representation of pushing the mower around the yard, but not a representation of taking each specific step. One might nonetheless judge that taking a specific step sufficiently approximates the general plan of pushing the mower around the yard.

Alternatively, one might feel strict about the approximation to the plan. Santiago Amaya appeals to the role intentions play in setting evaluative standards for action, and argues that the SPV fails to adequately explain the sense in which certain failures to achieve one's goals are really mistakes. Regarding cases like Mele's, he has this to say:

> Intending to achieve a certain goal normally involves performing some subsidiary actions. Insofar as one intends some of them as implementations of one's goals, these will be intentional… At the same time, it is often the case that one is perfectly indifferent about many of the subsidiary actions one performs in the pursuit one one's goal, even when those actions serve the purpose of helping one achieve that goal. These actions seem better described as non-intentional. As long as one manages to act as intended, not performing one of them would not count as an executive mistake. (Amaya 2018: 1784)

It's a merit of my account that it permits further specification, in terms of the SV or the SPV, along whatever lines the debate ultimately takes.

5.5.3 Senseless Movements

Some will wonder how this account squares up to senseless movements. Brian O'Shaughnessy (1973) saw in such actions—he lists "moving one's toes as one sits reading a book" as an example—a middle ground between intentional action and non-action:

> Does nothing lie between being a corpse-like graven image and a vehicle for reason? How else but as action is one to characterize the making of these movements, and to what but the person is one to attribute them? One can hardly telescope them into mere spasms on the part of the toes. And does one not suddenly become aware of doing them? Yet it is not as if they were intentionally senseless, as it were the 'small talk' of bodily movement, for they are not chosen . . . these trifling actions can express nothing more distinctive or more mental in our inner life than vague unease; and anyhow they are an afterthought in the scheme of things. They relate to standard examples of action somewhat as do objects that are mere lumps of stuff, say rough diamonds, to objects that are both lumps of stuff and more, e.g., artefacts, natural kinds. (But whereas all matter might have been in the form of mere chunks, these could not be the only examples of action in the universe.) Such senseless 'raw' acts are not amenable to interpretation, not even the interpretation of having no interpretation, and that is why they are not intentional under any description. Excluding them from the class of all actions would be roughly akin to excluding gold nuggets from the class of material objects.
>
> (O'Shaughnessy 1973: 366–7, fn. 2)

I reject O'Shaughnessy's claim that we can consider senseless movements neither intentional nor non-action. They are one or the other, and it will depend upon the case.[3] O'Shaughnessy hits on the right question when he asks whether we should attribute them to the person. Or rather, he is close

[3] John Hyman offers a longer list: "automatic reactions, such as ducking or drawing back one's head to avoid a blow, or making an involuntary adjustment to one's posture to maintain balance; some kinds of habitual action, including verbal tics such as echolalia (the automatic repetition of words and phrases spoken by a person one is conversing with) or interspersing speech with words or phrases like 'you know'; some kinds of uncontrolled action done in abnormal or pathological states of mind, such as panic or psychosis; unconscious action such as murmuring in one's sleep; and the spontaneous expression of emotion in facial expressions, vocalizations, and gestures, such as smiling, scowling, pouting, shrugging, and laughing, or crying out with pleasure or pain" (Hyman 2015: 50–1).

to the right question. I think the issue is whether such movements stem from central psychological states that we should attribute to the agent. Some idle movements qualify—as I write I twitch quite a bit, and it is possible to find in these twitches, I would think, a connection to urges to move based in vague unease or physical discomfort or some energy associated with the sentence at hand. Such behavior is intentional action, albeit a trifling sort. Other twitches and idle movements may be due to noise in my motor cortex, to reflex arcs further down the neurological hierarchy. There is less pressure to attribute these movements to me. They are then not action at all.

In agents like humans, who possess a complex neurological architecture supporting behavior, there may be in-between cases. We know that there are levels in a processing hierarchy supporting human action control (Fridland 2014; Shepherd 2015b; Christensen et al. 2016). We know that at times the agent's intentions can be at cross purposes with the somewhat independently structured processing supporting motor execution and sensorimotor adaptation (Day et al. 2016; Mylopoulos and Pacherie 2017; Shepherd 2019). It is an empirical question just how and just whether lower-level states like motor commands should be included within the agent's broader plan-state. Perhaps some should; perhaps some should not. I do not wish to commit here; the science is ongoing. The point is simply that if some movement follows in the right way from an agent's plan-state, and approximates some aspect of the plan sufficiently, we should think of it as intentional action. If the movement follows from something else—some state or some event that is outside the scope of any plan of the agent's—then it is not action at all.

5.5.4 Belief and Knowledge

My account of intentional action says nothing explicit about knowledge, and little about belief. One might, of course, think of the agent's plan-state or package of plan-states as involving or requiring belief. I remain neutral on that. One might also think of the execution of intentional action as necessarily involving knowledge how. On an anti-intellectualist view, on

Walking through each item on the list would be tedious. But I don't see any deeply problematic cases here. Some of our plan-states (i.e., ducking) are acquired very rapidly, and some involuntarily. Some cases tempt intuitions. But I do not think murmuring in one's sleep is an intentional action—it is an action-like effect of something one may be doing intentionally in one's sleep.

which knowledge how is an ability, I have no deep problem with that. Here I wish to discuss a different sort of view, regarding the importance of belief, and arguably knowledge that, for intentional action.

Intellectualism about knowledge how is the view that 'for one to know how to φ, for some task of φ-ing, is for one to know a proposition that truly answers the practical question "How could one φ?"' (Pavese 2016a: 650). Carlotta Pavese argues that knowledge how to A is necessary for possession of the ability to intentionally A (Pavese 2016a, 2016b, 2018).

Pavese argues as follows:

1] An agent J's success at intentionally A-ing requires that J have a true belief about how to A.[4]

2] Given 1], J cannot be in a position to intentionally A unless J has a true belief about how to A.

3] J cannot have the ability to intentionally A unless J is in a position to intentionally A.

C] The ability to intentionally A requires a true belief about how to A.

Ultimately, Pavese thinks that more than true belief is required for intentional success. Later in her (2018) she offers arguments based on cases as well as claims about what is required for satisfactory explanations of intentional success that conclude that intentional success at A-ing requires not just true belief, but knowledge of how to A. It follows that the ability to intentionally A requires knowledge how to A—a conclusion that affords the analysis described above. These latter arguments are interesting. Before discussing them, I discuss a potentially more basic issue.

It is clear that without premise 1], Pavese's argument will not go through. Pavese supports this premise initially by reference to the wide endorsement this premise receives amongst action theorists.[5] It is certainly true that many have endorsed a belief requirement on intending, and by extension on intentional action.

But what kind of belief is required? One way to get a belief requirement on intentional action more or less for free would be to endorse cognitivism

[4] As Pavese puts the thought, the intentionality aspect of the ability to intentionally A "reduces to a doxastic attitude" (Pavese 2018: 4).

[5] Among others, she cites Goldman (1970); Brand (1984); Harman (1976); Thalberg (1984); Ginet (1990); Mele (1992); Mele and Moser (1994), though we will see that Mele (1992) does not endorse it, nor, on my reading, does Mele and Moser (1994).

about intention. According to cognitivists, an intention just is a sort of belief. There is something to be said for this view (see Marûsić and Schwenkler 2018). But I will remain neutral on it for present purposes. In any case, it remains a controversial view in the philosophy of action, and Pavese explicitly states she is not here assuming cognitivism.[6]

The kind of belief requirement typically endorsed in the philosophy of action is motivated by the following consideration. An agent cannot be said to be genuinely committed to A-ing in the way an intention to A would require if the agent believes that she is likely to fail. So agents should believe that they are likely to succeed. The belief Pavese offers is in line with this: "If one successfully intentionally φs at t, then at t one believes, for some way ψ of φ-ing, that one is sufficiently likely to φ by ψ-ing" (2018: 6).

The viability of this belief requirement on intentional action depends on whether such a belief is really important for explaining intentional success. I am dubious that it is. Consider the following case due to Al Mele:

> In numerous instances of intentional A-ing, the question whether we will succeed in A-ing would appear to be the furthest thing from our minds. Yet in many such cases, the claim that we intend to A is unproblematic. A few minutes ago, while I was typing, I heard a knock at my office door. As is my habit, I answered the knock. I answered it intentionally, and the suggestion that I intended to answer it would encounter little resistance. Yet I do not remember having a belief at the time to the effect that I (probably) would answer the knock. Indeed, I seem to recall that I had no such conscious belief. Moreover, given that my intention . . . can do its work without the assistance of a belief—conscious or unconscious—that I (probably) would answer the door, there is no apparent need to postulate the existence of this belief. (Mele 1992: 147)

Kieran Setiya offers a similar case and verdict:

> I need no more than the intention to clench my fist, and the disposition to do so under the guidance of that intention, in order to clench my fist intentionally as a basic action . . . What [the intellectualist] owes us . . . is an account of what it is to have the relevant intention and disposition on which they entail the belief that I could clench my fist by intending to do

[6] Could one use cognitivism to get to an interesting conclusion regarding the relationship between knowledge how to A and the ability to intentionally A? Probably so, in my view.

> so. The problem here is not simply one of omission, that there is more to say, but that it is difficult to see how any plausible account of intention or guidance would entail this belief. (Setiya 2012: 294)

Such cases motivate a negative belief criteria on intending (and thus, on intentional action): "J intends (at time t) to A (at time t') only if he does not believe (at t) that he (probably) will not A (at t')" (Mele 1992: 147). (We could expand this to plan-states generally.) This criteria would spell trouble for the first (and a key) premise in Pavese's argument. Should we accept it? Beyond the intuitive nature of the cases, the primary consideration Mele and Setiya advance is that the belief is doing no relevant work in explaining the intentional success.

Pavese, however, advances arguments in favor of the explanatory relevance of a belief requirement. I will consider two. The first is offered in response to Setiya. Pavese distinguishes between two notions of intentionality. Intentionality-minus is "a property that an action-token possesses just in case (1) the action is intended and (2) it is caused in the right non-deviant way by that intention" (Pavese 2018: 11). Intentionality-plus adds to this a belief requirement. Pavese then argues that basic actions should only qualify as intentional if they meet intentionality-plus. Her reasoning is as follows:

> Basic actions for a subject at a time must be actions a subject is, at that time, able to plan from, if only the subject were to engage in a little [bit] of reflective effort. But in a circumstance where I explicitly (that is, upon considering the question) did not believe that I can perform an action (such as clenching my fist) by merely intending to, I would not be in position to plan a complex action that has clenching my fist as its part.
>
> (Pavese 2018: 11–12)

There is something to be said for Pavese's claim about being able to plan from actions. But in any case this reasoning does not work. It is consistent with a negative belief requirement on intentional action that basic actions be available for planning. So the claim does not uniquely support a positive belief requirement.

The second argument of Pavese's I wish to consider she calls the argument from verbal feedback. The argument involves two moves. First, we acquire knowledge how through verbal feedback and through trial and error. Second, if we did not possess standing belief states about how to A

that verbal feedback manipulated, then explaining how verbal feedback works would be difficult. Pavese writes:

> Suppose my tennis teacher tells me that I am not holding the racket correctly. It seems natural that their feedback will result in revising my beliefs about how to hold the racket. More generally, revising one's belief about how to perform an action seems a natural way in which know-how can be acquired and improved through verbal feedback. If Belief requirement on know-how is correct, this aspect of know-how can be easily captured. By contrast, if know-how did not require a standing belief state from one occasion to another, the acquisition of know-how through verbal feedback and trial and error could not be a matter of revising one's belief.
>
> (Pavese 2018: 8)

In response, we can grant that the acquisition of knowledge how can be influenced by changes in beliefs. This may happen via many mechanisms, including the changing of the trial and error process that may be more critical for knowledge how acquisition. But the influence of belief on knowledge how acquisition—or on acquisition of new abilities to intentionally A—does not entail the necessity of belief for such acquisition in every case. If there are cases where know how, or the ability to intentionally do a new thing, are acquired without verbal feedback, this places the argument from verbal feedback in doubt. And it does seem like there are many such cases. When a young Björn Borg spends hours volleying a tennis ball to himself off of a garage door, he seems to be building his base of abilities to intentionally hit the ball in various ways. But there is no verbal feedback offered.

The main point of contention here concerns the relevance of beliefs about how to do something to the explanation of intentional successes. I am claiming beliefs are not necessary. A different way to test this claim may be useful for people whose intuitions are stuck in the middle.

Let us consider a case of contrastive explanation of intentional success. The case of Dan and Dave:

Dan and Dave are sprinters who have trained together since childhood. Their practices are identical, and in learning their technique their coach is careful to teach them the exact same methods—for he believes his methods are uniquely the best for successful sprinting. Over time Dan and Dave internalize these techniques so well that they can coach each other, which

they do by reminding one another of the same mantras, and by stressing the appropriate ways to conduct each training exercise. Indeed, by the time their coach is ready to allow Dan and Dave to compete, the propositions that they explicitly believe about sprinting and how to do it are exactly the same. Now, as it happens, Dan is slightly faster than Dave, although as their coach does not allow them to race each other, neither knows this. Indeed, neither knows just how fast they are—for their coach does not release their times. But Dan is able to reliably run sub-10-second 100-meter dashes, while Dave cannot break 10 seconds. When Dan and Dave run their sprints, they intend only to run fast, utilizing the techniques they have learned.

Here is a verdict on this case that I find intuitive. When Dan runs with the intention to run a fast 100-meter dash, he not only intentionally runs fast, he intentionally runs a sub-10-second 100-meter dash (9.92 or whatever). (Accepting this intuition pushes one to the SPV, but that is hardly a cost.) When Dave runs with the same type of intention, he not only intentionally runs fast, he intentionally runs a near-10-second 100-meter dash (10.06 or whatever).

Now, both Dan and Dave may or may not have beliefs about their likelihood of success. This case is not an illustration of the negative belief requirement on intending or intentional action. The point is that Dan and Dave have different abilities. Dan is able to intentionally run sub-10-second 100-meter dashes. Dave is not. But their sprinting-relevant beliefs are conceivably exactly the same. So their beliefs are doing no work in the explanation of the difference in their abilities, and hence of the difference in their intentional successes. Very plausibly, what is doing the explanatory work here are differences in physicality.

One might reply that Dan does not have the ability to intentionally run sub-10-second dashes, since Dan is not intentionally doing so in the case I described. That violates intuition—after all, Dan is in control of his behavior throughout the run, and he runs a sub-10-second dash. Resistance to this verdict must be based in some ancillary theory. Perhaps on the basis of a principle like this: "An agent intentionally A-s only if she is aware of A-ing." Dan is not aware that he is running a sub-10-second dash. But, as others have pointed out, this principle is problematic. Mele (2001) offers a case in which he intends to make his daughter laugh by sending her an e-mail. He is not aware that this is what he is doing as she opens the e-mail and laughs. But he intentionally makes her laugh all the same.

Pavese might appeal to her theory of knowledge how to reply that Dan and Dave are different in an epistemically relevant way. In other work

Pavese has argued that when an agent knows how to A, she possesses a practical concept that is a component of a proposition regarding how to A, the grasping of which enables her to follow a rule to A (Pavese 2015a, 2015b). One claim she defends along the way is that knowing a rule to A is necessary for knowing how to A (Pavese 2015b). Perhaps, then, Dan knows a rule that Dave does not. Dan knows a rule for running sub-10-second dashes. And since rules are components in crucial propositions, perhaps Dan's knowledge how to run is distinct from Dave's in a way friendly to the intellectualist about knowledge how.

But we can be more explicit about the case, telling the story in such a way that the rules Dan and Dave come to know are exactly the same, and that the only difference between them has to do with physical features such as the elasticity of their calf muscles. The intellectualist about knowledge how will not want these features to make a difference to their knowledge how. Indeed, intellectualists are explicit about ruling such cases out (see, e.g., Pavese 2017: 376). So the intellectualist will need to deny that Dan and Dave have different abilities to intentionally act. But of course Dan and Dave do have different abilities to intentionally act. Dan can intentionally run sub-10-second dashes, and Dave cannot. Once their coach lets them race on the international stage, this important difference in their abilities will be a part of the reason that Dan becomes famous, and widely acclaimed as a great sprinter, while Dave does not.

This case highlights an additional way that beliefs are superfluous to the explanation of intentional successes. In conjunction with cases like Mele's and Setiya's, it should lead us to reject such a requirement. The result is that Pavese's argument, leading to the conclusion that the ability to intentionally A requires intellectualist knowledge how to A, should be rejected.[7]

5.5.4.1 Explaining Intentional Success

I want to discuss Pavese's arguments regarding the explanation of intentional success. The arguments are intrinsically interesting. But my rationale is to highlight an important difference in approaches to action like Pavese's and those like mine.

[7] A further result follows from the conjunction of the above considerations with the principle elucidated at this chapter's beginning. Pavese accepts that if an agent can intentionally A, then that agent knows how to A. As we have seen, Dan can intentionally do something Dave cannot. So Dan knows how to do something Dave cannot. But Dan and Dave can be described in a way such that their beliefs are the same. So the intellectualist about knowledge how has to reject this principle in order to hold onto intellectualism.

Pavese first argues that there must exist a safety constraint on the belief(s) that figure in knowledge how. Weak safety holds when a belief is true "in most (or sufficiently many) of close cases." Safety holds when a belief is true in all the close (i.e., similar) cases. She links ideas thus: "only modally robust beliefs can figure in satisfactory explanations of success and . . . know-how figures essentially in explanations of successes" (Pavese 2018: 14). The reason is that an explanation of a success needs to show that the success was not a mere fluke—was not an accident, was not deviant, was not due to too much luck.

Pavese uses an example of Daniel Greco's (2016) to illustrate the belief case. Alice makes it to the Colosseum. The explanation may very well refer to true beliefs about the Colosseum's location. Pavese says that "Alice cannot intentionally have ended up at the Colosseum without true beliefs about how to get there" (16). And these beliefs should be safe—an accidentally true belief will not show that the success was not deviant.

Regarding this explanatory principle that an explanation of a success needs to show that the success was not a mere fluke, we are in agreement. However, belief does not always play a role in explanations of intentional action—sometimes, intentions or other plan-states will do just fine.

And there is an interesting difference that emerges here. (A version of this difference will re-emerge in chapter seven, when I contrast my account of skill with accounts like Pavese's.) The difference is that the aim of a plan-state (like an intention) is different from that of belief. A plan need not be perfectly veridical to be good. The account of intentional action I have offered requires that the plan be a good plan—the agent's exercising control with respect to the plan or part of the plan is a reliable way for the agent to meet with success—a success rightly deemed an intentional action. It is possible, even if unlikely, for a plan to be good—for a plan to be safe or modally robust regarding the connection to success—without embedding truths about the environment in which an agent acts.

Harriet is a high jumper with a perceptual defect. She systematically perceives, and comes to believe, that the bar is higher than it is. This leads her to form plans to jump with a certain amount of force that reliably leads to her clearing the bar. Suppose that were it not for this perceptual defect, she would be less reliable at clearing the bar. Here is it a false belief that helps explain her success.

Intentional action's proximity to success sometimes crowds out its proximity to truth. This furnishes a further reason why intentional action does not require knowledge.

5.6 A Problematic Case (Systematic Gettierization)

An interesting problem case can be drawn from the literature on Gettier cases. The relevance of Gettier cases to action runs via structural similarities: it has been suggested that, given parallels between knowledge and action, such cases should apply to the case of action (Sosa 2015; Miracchi 2017).

The kind of case I have in mind is what Miracchi (2015) calls a systematic Gettier case. The structure of the case is such that an agent engages in a process, or exercises a competence, that usually produces some product, be it knowledge-formation or action-production. But in this case the conductance of the process or the exercise of the competence is unusual. First, the agent undergoes what Miracchi calls bad luck. Something happens that, in other circumstances, would completely derail the knowledge-formation or action-production. Second, however, the agent encounters what Miracchi calls good luck. Some feature of the situation or the process corrects for the bad luck, such that the agent forms a belief, or produces behavior, that accords with the expected or normal result. What's more, it looks like the agent's performance is explained by her conductance of the process, or her exercise of her competence. And yet the belief is claimed not to amount to knowledge; the behavior is claimed not to amount to intentional action. For the luck involved renders the success a result of deviant causation.

To get a better feel for the structure of the cases, and to see the parallels from knowledge to action, consider the following two cases. The first is taken from Miracchi (2015: 39), and the second is inspired by a case Sosa (2015: 13–14) discusses:

DOUBLE TROUBLE: Creola's two friends Fred and George sometimes like to play tricks. Fred goes on a month-long trip, and communicates with George daily. Fred decides that he will play a trick on George and tell him the complete opposite of what actually happens to him on his trip. George, charged with relaying news of Fred to Creola, decides he will play a trick on her, and tells her the complete opposite of what Fred tells him. As it so happens, Fred and George's subterfuges reliably cancel each other out, so that Creola reliably receives true information about Fred's trip. She forms beliefs on the basis of George's testimony just as she would in any other standard case of belief-formation on the basis of testimony.

THE SHOT: Robin notches an arrow in his bow and prepares to fire at a far-off target. He knows where the target is, because his friend John told

him, but he does not inspect the scene in any detail. This is normal for Robin: he is an accomplished archer, and he and John enjoy this game in which John tells him of a target, and Robin hits it. Robin intends to shoot a shot he is confident will hit the target, even though the shot must curve around a bend, out of sight, in order to do so. Unbeknownst to Robin a strange wind swirls around the bend. The arrow will pass through it, thereby deviating from the target. But the deviation is such that the arrow's tail will graze a nearby tree, putting the arrow back on target.

We are told to think that Creola forms a true belief by exercising her normal competence to believe truly on the basis of testimony:

> We may suppose that, had Creola not formed beliefs using her default competence to believe truly on the basis of testimony, she would have attended to subtle signs in George's voice, eyes, etc., that would reveal his subversion. She thus would have believed he was lying, and have formed false beliefs about Fred's trip. Thus her believing truly, and not just her believing, is causally explained by her competence. (Miracchi 2015: 39)

And yet, we are supposed to think, Creola's belief does not amount to knowledge. So the exercise of a reliable competence to believe truly is not enough for knowledge.

What about the structurally similar case regarding action? Robin executes an intention as he had planned. And his behavior was, in the specific circumstances at hand, a reliable way of satisfying the intention—of hitting the target. And yet one might have the intuition that Robin did not hit the target intentionally.

Ernest Sosa shares that intuition, but pushes back nonetheless. His account of non-deviant causation is similar to my own, but differs on an important detail. Given this—and given that his view of knowledge is a primary target of the systematic Gettier cases in the first place—his response is worth considering.

5.6.1 Sosa's Response and my Own

Sosa's response is situated in a broader competence-based account of intentional action. This account is of a piece with related competence-based accounts of knowledge and perception. But the focus here is action.

According to Sosa, in intentional action the agent must not only reach her aim, she must do so by way of a competent performance. Sosa claims that "aptness—success that manifests competence—is the key to 'the right way'" (Sosa 2015: 19).

It is important for Sosa that the agent's competence possess a certain structure. Sosa calls this a SSS structure. The agent must have Skill, which for Sosa is a basic package of dispositions that enable successful performance. The agent must also have Shape, which is to be in the right circumstances. For example, if the agent is driving, she should be awake and sober and so on. To Skill and Shape, Sosa adds Situation. So, regarding driving, Sosa states that the car should actually work, the roads should be passable, and so on.

This is clearly similar to my own understanding of non-deviant causation. But there is an important difference. Although Sosa embeds skill within shape within situation in the good case, they appear to be separate. That is, an agent may possess the skill to drive even when she is not in the right circumstances, either because the Shape or the Situation are wrong. This hearkens to a distinction between general and specific abilities (Mele 2003a). I do not oppose the usefulness of this distinction for some purposes, but here it may mislead. For I claim that the agent does not possess what Sosa calls Skill in the absence of any circumstances whatsoever—some set of circumstances must be posited in order to get the dispositional structure off the ground. It is better, especially when dealing with difficult cases like the present set, to be explicit about this.

Sosa is not explicit about this, and as a result his response to the cases runs in a different direction than my own. Sosa's verdict is that in cases like THE SHOT the agent does not intentionally hit the target. Sosa requires three features for intentional action: success, competent performance, and that the success itself manifests competence (2015: 19). Robin's performance has the first two, but not, Sosa avers, the third. Here is what he says of a case similar to THE SHOT:

> Why is this shot not apt after all? A performance is apt when it succeeds because of the agent's competence. But our archer's wind-aided shot does seem to succeed because of his competence! If the agent's competence had not resulted in the right orientation and speed upon release from the bow, then the arrow would not have hit the target.
>
> Taking a leaf from Davidson and Grice, we might judge success to be apt only if it derives causally from competence in the right way. Success

> essentially aided by lucky gusts of wind would not derive in the right way from the archer's competence. (Sosa 2015: 13–14)

Sosa is aware that this seems to introduce circularity into the analysis. Non-deviant causation is causation in the right way. If all we can say is that this amounts to success that derives from competence in the right way, we seem to have said too little. But Sosa argues he can go beyond this appeal to the right way. He does so by appeal to a primitive relation of manifestation.

An agent's success derives from a competence in the right way when it manifests that competence.

Manifestation is primitive. We can say little about it explicitly. We grasp it via intuition.

Sosa offers an analogy with a case regarding dispositions. A glass is fragile because it is disposed to break when struck. But suppose the glass, while about to smash on the floor, is instead zapped by a powerful ray that would shatter even an iron dumbbell. Sosa claims that, intuitively, the fragility of the glass is not manifested in the right way. And this kind of intuition is enough to justify claims about non-deviant causation in the relevant cases. So we can say, without offering any further explanation, that in some cases Robin's success will manifest competence, and in some cases not. The differences are picked out by intuition, and these discriminate deviant from non-deviant causation.

This is a bold line to take, and one with a certain attraction. Cases of deviant causation are, after all, picked out by intuition. Why not cases of non-deviant causation?

But relatively brute appeals to intuition can easily mislead. So, what of the intuition that Robin does not intentionally hit the target in THE SHOT? Here intuitions mislead. THE SHOT, and similarly structured cases, tempt us to imagine a different set of circumstances than the one we are explicitly told to consider. This different set of circumstances involves unreliability in the process. Fred could easily err, and tell George the wrong thing. The wind in Robin's case could easily swirl a different way. When discussing these kinds of cases Miracchi uses the term luck, explicitly suggesting that there is something odd about the relevant processes.

So the oddity is just that in the circumstances we are tempted to imagine, the processes would prove unreliable. But these are not supposed to be simple deviance cases. The Gettierization is supposed to be systematic.

But who will stand surety for the modal guarantee? It is the selector of the systematically Gettiered set of circumstances.

And how then do we imagine a case in which the systematically Gettiered process is unsafe? The problem is to be found in one of two places.

In some cases, the systematic Gettierization will be complete, and the problem will be with our imagination. We will wrongly imagine a simple deviance case. And we will wrongly attribute luck to the situation. In this case, however, in the relevant circumstances, these processes are not lucky or unlucky. They are a part of the fabric of things. They are how things work.

So consider a case where things always work in the way they work for Robin the archer. True, Robin cannot intervene in the part of the process that involves the wind and the tree. But we have already considered a requirement of intervenability on control. The lesson here is the same as with the dice. We reject such a requirement. Many intentional actions have significant components that operate ballistically or opaquely.

In other cases, the problem will be with the set the selector has constructed. If things do not really consistently work that way in the type of circumstances specified, then we do have a case of luck, but the luck is no problem. The systematic Gettierization is not complete—it is rather an artefact of the subset of circumstances under specification. And if so, the luck undermines the behavior's qualification as intentional action, and not the account of intentional action under consideration.

As I said above, there is nothing magic about non-deviant causation. It does not matter how ballistic or opaque parts of causal processes operate, so long as they are sufficiently reliable. Systematic Gettier cases have such processes, and agents in these cases act intentionally.[8]

5.7 Causalism and Reductionism

I offer an account of intentional action within the causalist tradition. The non-causalist's most common complaint against the tradition is deviant causation. Underlying this complaint, however, is a potentially deeper issue. For deviant causation is usually mentioned because the non-causalist finds something malignant in what she reckons the more general causalist outlook. The worry is that causalism has reductive ambitions, and that these ambitions are out of place in an account of intentional action. (Why this is so will differ depending on the non-causalist.)

[8] There may be lessons here for those seeking parallels between intentional action and knowledge.

Perhaps some causalists have a reductive ambition. But perhaps they hold this ambition in the way that a philosophy lecturer holds an ambition to amuse her students. Being funny is a thing additional to the core aims of the philosophy lecturer (or assume it is for illustrative purposes). She may succeed in her ancillary aims or not, and her success or failure is apart from her success or failure in teaching philosophy. So too, I suggest, with causalism and reduction. Qua causalist, we need only commit to the existence of a causal condition on the nature and explanation of intentional action.[9] And it is not difficult to see the attractions of this minimal commitment. For intentional action is about things that agents do, things agents make happen, things agents bring about. Some think the causal relation should hold between psychological states and behavior. Some think it should hold between reasons and behavior. Some think it should hold between agents and behavior. Some think it should hold between facts regarding one of these things and behavior. Be all of that as it may—to claim that a causal condition is not necessary for an account of intentional action is to allow intentional actions that take place in spite of the absence of a causal relation between whatever relata one prefers. It is to allow intentional actions that neither the agent nor any relevant agent-involving process causes in any way.[10] That strains credulity.

What, then, are non-causalists upset about? That may depend. Here I discuss two ways non-causalists have pressed causalists regarding reductionism. In each case the worry is pressed as though it is essential to the causalist view. That is in part why I feel compelled to address it. I do not see the reductive ambition in the way many non-causalists seem to. Perhaps, then, there is some prospect for peace. But we will see that there

[9] Something like the following proposal, due to Mele (2017b: 30–1), should do:

> D5. Necessarily, if E is an adequate explanation of an intentional action A performed by an individual agent S, then E cites (1) a reason that was a cause of A or (2) a belief, desire, or intention that was a cause of A or (3) a neural realizer of a belief, desire, or intention, which neural realizer was a cause of A or (4) a fact about something the agent believed, desired, or intended, which fact was a cause of A.

[10] One might say instead that all the non-causalist is commited to is a claim about action explanations: action explanations need not reference any causal condition (see Wilson 1989; Sehon 2016). Or, as Sandis usefully glosses the view, "it is a sufficient condition of a reason-giving explanation of action that the reason cited renders the action intelligible" (Sandis 2006: 12)—where the reason cited need not reference causation, even obliquely, in order to do its work. But it is difficult to avoid the implication that, if causation plays no role in the explanation of intentional action, then it is possible to have actions that do not involve causation. See Mele (2003b, 2019) for discussion.

is a sense in which the perspective I endorse is reductive. So the peace can be, at best, partial.

5.7.1 Basic Action and Intentionality

Douglas Lavin has pressed a worry against causalism's reductionist ambitions by way of a discussion of basic action. Roughly, basic action is an action an agent performs without performing any other action. Lavin argues that causalists need a notion of basic action. According to Lavin the legitimacy of "the seemingly innocent idea of basic action . . . is vital to the intelligibility of the causal theory of action" (Lavin 2013: 274). Why so? Lavin explains:

> Its legitimacy is vital to the intelligibility of the causal theory of action, according to which physical action consists of a mere event and a condition of mind joined (in the right way) by the bond of causality. Left unchecked, means-end reason threatens to permeate physical action, and thus threatens the sovereignty of the sphere of material events at the center of the causal theory: such events, including the movements of one's body when one intentionally moves it, are thought to be constitutively independent of the subject's rational capacities. Basic action is a necessary countermeasure, a sort of metaphysical containment wall needed to preserve the separate jurisdictions of the mind of the acting subject and what merely happens. (Lavin 2013: 274)

Lavin is saying a lot. Obviously, he is working with a certain interpretation of the causal theory, and the metaphysics to which the causal theory is committed. I do not doubt some causalists would recognize their view in Lavin's depiction. I do not recognize my own, however.

Like some (but not all) causal theorists, I have in the past been happy to think of basic action as a useful but inessential notion. As Jennifer Hornsby puts a thought I am happy endorsing, "there must be something right in saying that no one would do anything if everything she might do was something she could only do by doing something else" (Hornsby 2013: 2). Before reading Lavin's paper, I would not have thought that my endorsement of Hornsby here was an endorsement of anything vital to the causal theory, nor of anything resembling a metaphysical containment wall separating the

mind from what merely happens. If Lavin is right, however, that is exactly what I am endorsing. Why does Lavin think this?

Lavin thinks that without basic action, causalism cannot give an account of the structure of intentional action. Consider two representative passages:

> Consider illuminating a room, building a house, or baking a cake. How do we execute complex projects such as these? The answer, says contemporary action theory, is that we perform complex actions by performing a more or less intricate sequence of basic actions, while we perform basic actions immediately, directly or 'just like that'. (Lavin 2013: 273)

> The classifications [of actions as basic and non-basic] is meant to be one we must recognize if we are to understand the very structure of intentionally doing something: whatever large-scale projects one has realized through the ordering of means to ends, one must eventually reach a fine enough resolution and come upon things that have been done without any thought about how to get them done. (Lavin 2013: 276)

Here Lavin gives two characterizations of his target. Basic action is action performed "just like that," and basic action is action performed with no thought about how to get it done. Lavin clearly thinks that the existence of such actions is necessary for a causal account of the structure of intentional action. Unfortunately, without further explication, it is not transparent what Lavin means with notions such as actions performed "just like that" and without thought about how to get them done. Fortunately, Lavin offers a more explicit explication of what he has in mind in the following passage.

> [Basic action can be described] through the concept of an end: a basic action is not the end of any other action; nothing else is done in order to do it; it is not an answer to 'Why?' when asked about any other action. And equally through the concept of a means: no means are taken in the execution of a basic action; it is not done by doing anything else; there is no answer to 'How?' when asked of it. (Lavin 2013: 275)

I will eventually suggest that all these descriptions of a basic action, which are intended to amount to the same thing, in fact do not. Before I get there, however, I want to diagnose why Lavin thinks basic action is needed as a metaphysical containment wall without which causalism would fail to achieve its central motivation. I want to do this because I agree with Lavin

that basic action is dispensable, but I disagree that this makes any problem for causalism.

Lavin thinks basic action is essential to causalism because he sees causalism in a certain way. Lavin thinks causalism is committed to the existence of actions without means-end structure. Basic actions are meant to be practical atoms out of which non-basic actions—those with a means-end structure—are built.

What happens, then, if we show this notion of basic action to be unintelligible? The key aim of causalism is revealed as a failure:

> In giving up on basic action, the means-end order is shown to be at once an order of causality (the means realize the end) and an order of reason (the end rationalizes the means). And this order, an order of practical reason, is shown to be internal to what happens, to the progress of the deed itself. (Lavin 2013: 274)

Or:

> Without basic action . . . a residue of intentionality would always remain in the representation of material processes themselves, of movements as movements. It would not be possible to realize the explanatory ambition of the causal theory, namely to fit action into a world that does not contain intrinsically intentional material processes, unless basic action is the fundamental manifestation of rational agency. (Lavin 2013: 279)

According to Lavin, if causalism cannot remove the residue of intentionality in the representation of material processes, it fails. Similarly, if causalism cannot get rid of the rationalizing of means via ends—presumably, by explaining away such rationalizing in causal terms—it fails. Causalism cannot do these things, Lavin argues, because it relies on basic action to do these things for it.

Suppose we did remove the residue of intentionality from the representation of material processes. What would we have done? On a broad reading of "removing the residue," we might have just naturalized intentionality, where to naturalize intentionality is to show that the semantic (and/or intentional) is not "permanently recalcitrant to integration in the natural order" (Fodor 1984: 232). How that should go is a topic for another place. Here we want to know: is the naturalization of intentionality something to

which causalists must be committed, qua causalists? Certainly many causalists have wanted to fit a workable notion of intentional action into the natural order. But one way to do that is to show that there is a workable notion of intentional action that—while not itself naturalizing intentionality—poses no special problem for the naturalization program.

Perhaps Lavin only means to give causalism the task of naturalizing action theory. But if so, it is not clear why it is a problem that a residue of intentionality remains. It this connection, it is instructive to consider how one of the great causalists, Myles Brand—in a book with the subtitle *Towards a Naturalized Action Theory*—understood the causal theory. In his (1984) book's first chapter, Brand explicitly contrasted the causal theory with what he called the Oldtime Volitional Theory. The latter theory is reductive, aiming to show that "human action consists in causally related nonactional events" (Brand 1984: 7). Brand rejects this theory, noting that he sees no good way it can avoid essential use of action terms in its definition of action. Instead, he proposes the causal theory, which "does not attempt to provide a reductionist account, which it seems clear cannot be sustained" (Brand 1984: 17). Brand is trying to naturalize action theory in one sense, namely by preparing an action theory that is amenable to, and that can be integrated with, a scientific study of human action. But Brand—qua causalist—is not trying to remove the residue of intentionality from his account of action.

I think this is how many causalists view one motivation for causalism. Causalism, as a theory of action, is free to be agnostic over the prospects for naturalizing intentionality.[11] And if so, the causalist need not worry over the residue of intentionality in her representation of material processes, any more than the cognitive scientist does. Whether the residue can be removed depends on progress in naturalizing intentionality much more broadly. Qua action theorist, the causalist is free to utilize intentional mental states (e.g., intentions) and processes (e.g., intentional activity).

How does causalism look without the kind of practical atomism that bothers Lavin? Recall that one way Lavin describes basic action is in terms of means: "no means are taken in the execution of a basic action; it is not done by doing anything else; there is no answer to 'How?' when asked of it" (Lavin 2013: 275). Consider two kinds of actions. For the first kind, illuminating answers to how questions are available. How do I bike from Catherine

[11] Causalism may—depending on the brand of causalism in play—remove the existence of any special action-theoretic problem for the naturalization program.

Street to the Chester Arms? I get on my bike, start pedaling, turn left on Magdalen, right on Stanley, and so on. But there are also actions for which I cannot offer illuminating answers. How do I twitch my thumb just so? I have no explicit knowledge of how I do it, other than to say I know how to form an intention to do it and I know how to direct effort towards the satisfaction of such an intention.

(Presumably Lavin would deny that my directing effort towards the satisfaction of such an intention counts as taking means, but I'm not sure why. Basic actions can be effortful, and we can deliberate about whether the effort is worth it. An agent may think: "Perhaps the means taken to achieve a different end would be preferable." Such an agent does not seem deluded.)

We should not take my own limitations with respect to thumb twitching to signal a limit on what can be said regarding how the thing is done. In principle, for any segment of activity B within some action A, there could be an agent who could provide an illuminating answer regarding how B is done. Perhaps such an agent would need to possess an infinite practical intelligence and bodily abilities to match. Supposing such an agent conceivable, this agent would always be able to break down components of her A-ing into further components. The agent could go on for an infinite amount of time, telling you how the infinitely small subcomponents of her action are executed.[12]

[12] Does this mean that there must be an infinite series of intentional actions the agent must perform in order to intentionally A? Michael Thompson (2008), imagines a case in which one pushes an object from α to ω, and in doing so pushes the object past any of an infinite number of places. Thompson thinks the following:

> It seems that we must also allow that I am doing each of these things intentionally, and moreover that I am doing each 'because' I'm doing the next one. And so, even though the imagined series of isolated positions has an obvious geometrical limit in ω, it seems that an interlocutor and I might together forge a potentially infinite sequence of perfectly legitimate questions and answers, 'Why?'
>
> (Thompson 2008: 113)

I don't deny this is possible in principle for certain kinds of agents. But as applied to human agents, here is a reason to demur. I have offered an account of intentional action on which some bit of behavior B is an intentional action only if B sufficiently approximates the representational content of some relevant plan-state. Now, suppose it is possible for there to be some agent capable of representing not only some prospective action A, but also an infinite number of segments of behavior that together comprise A-ing. If so, then it might be possible for there to be an infinite series of intentional actions an agent will perform on her way to intentionally A-ing. But such an agent will not be a human agent. Human agents do not represent their actions in such ways. So, suppose our infinitely intelligent agent acquires an intention to twitch her thumb that is coded in the same kind of formats as a human intention. In discussing her action of thumb-twitching, her descriptions of her own behavior will at some point cease to sufficiently approximate the content of her intention. As a result, she will no longer be describing anything she does intentionally.

Thompson has this to say:

The causalist need not be committed to the existence of practical atoms, nor to the existence of actions the agent performs without taking means. This means that, for all the causalist says, a residue of intentionality remains. But the causalist—qua causalist—should never have doubted this (if she ever did).[13]

Given that basic action is dispensable, why does basic action of a sort seem inevitable? And why does it look like a useful notion? Why does Hornsby's thought—"there must be something right in saying that no one would do anything if everything she might do was something she could only do by doing something else" (Hornsby 2013: 2)—look so plausible?

My answer is inspired by Hornsby's recent discussion. She outlines a kind of basic action according to which an action A is basic for an agent if she can intentionally A and if she lacks means-end knowledge of how to A (Hornsby 2013: 16–17). Depending on how one understands knowledge of how to A, this may not be quite right. What Hornsby seems to have in mind here is a lack of explicit means-end knowledge of how to A—a lack of the kind of knowledge that would allow an agent, in the course of a piece of practical reasoning, to put segments of her performance into the form of a practical syllogism.[14] It is certainly true that agents sometimes lack explicit knowledge of this form. So I think Hornsby is onto something important.

> Of course, I might put an end to this torture at any one of the interpolated points, saying, 'Well, I'm pushing it to ϕ, you know, because I'm pushing it to ω.' But this doesn't show that any of the intervening 'because'-statements that I have thus left unframed would not have been perfectly legitimate and true. (Thompson 2008: 113)

That may be true in a sense. But there is less reason to think that, for example, my pushing the object from $n_{4238948743}$ to $n_{4238948744}$ is an intentional action. It may be no part of my plan, even if my plan entails such a movement. Might we nonetheless describe the movement as an intentional action? I suppose that we could try, although many would find this under-motivated. My current point is that there is good reason to deny that we must describe the movement as an intentional action.

[13] In fact I suspect many action theorists do not have a fundamental need for basic actions. Two of the better book-length examples of causalist theorizing—Brand (1984) and Mele (1992)—make almost no mention of basic action. And one causal theorist Lavin cites as requiring basic action actually admits that one could build an action theory without reference to basic action (Enç 2003: 47). It is true that Enç goes on to build an account of basic action and make it central to his causal theory. But in my view this is because he thinks that the existence of basic action is an empirical fact about how human action is constructed, and so, as Enç writes, "a definition of action that acknowledges them would seem to be more informative" (Enç 2003: 48). This suggests that Enç is less interested in a conceptual question about the nature of action and more interested in human action, and that he conflates the two issues.

[14] Enç (2003) seems to have something similar in mind. Compare also Enç's definition of basic action, which appeals to an agent's bringing about some result R "without using her knowledge of how to bring about any other event in order to bring about the *result*" R (Enç 2003: 75). By "use of knowledge how," Enç appears to mean use in explicit practical reasoning.

I think basic action seems inevitable to us because we are not agents of infinite practical intelligence. There are bits of behavior that we can perform, even though we have no explicit knowledge of how we do such things beyond the relatively unhelpful knowledge that we do them by trying.[15] Some mental actions seem almost essentially this way. How do I remember things? Sometimes I go through subsidiary mental actions of rehearsal. Other times, however, I simply direct effort towards remembering, and somehow I remember. The same is true of some physical actions. How do I twitch my thumb just so? As I said above, I have no explicit knowledge of how I do it, other than to say I know how to form an intention to do it and I know how to direct effort towards the satisfaction of such an intention.

This, I claim, is a useful notion of basic action even though it commits no one to anything like a metaphysical containment wall needed to preserve a jurisdiction between my mind and what merely happens. It is useful because it marks out an interesting feature of human agency—our knowledge how to do things often outruns our explicit knowledge of how they are done.

5.7.2 Action First

Recently an approach to the nature of action has been gaining some traction. Yair Levy (2013) calls this approach intentional action first, a name that is meant to echo Timothy Williamson's knowledge first approach in epistemology. Lucy O'Brien writes instead of action as prime. In both cases the motivation is the same. We are to give up on offering a reductive analysis of action. Instead, we take action as metaphysically basic. Once we have done so, we can then use the notion of action to explain other items in the philosophy of agency.

The rationale for giving up an analysis of action is thus two-fold. There is a positive side—the notion of action is explanatorily useful, and we need it. There is a negative side—attempts to explain action, or to reductively analyze it at least, have failed. As Levy writes, citing the deviant causation literature "the continued failure to vindicate [a causalist] analysis merits exploring alternative, arguably more promising, research programmes" (2013: 710). O'Brien's opening salvo is rather direct:

[15] There are actions, in my view, which we do not explicitly know how to try to do, though we are able to try to do them, and though we are able to try to do them intentionally.

> Philosophers of action very often start with the question: what happens when someone acts? . . . I am going to go on to urge that there is a certain kind of answer to that question that is often expected, but that we cannot have, and that we do not need. (O'Brien 2017: 265)

O'Brien's position is useful in that, among its other merits, it is clear that her problem is not with the involvement of causation in action, but rather with a reductive ambition in accounting for action. O'Brien lists and discusses a range of potential necessary conditions upon action. Roughly, these can be conjoined into the claim that when I act "I, myself, change in a way that is up to me" (O'Brien 2017: 268). O'Brien's problem is with any account of action that would take a list of necessary conditions—whether hers, or some further analysis of hers, or some alternate list—and claim that they are jointly sufficient in a way that provides an analysis. This would be to "attempt to analyze my action in other terms" (2017: 270). This is to be resisted. For this "would imply that the act—my raising of my arm—is not actually a single unified element in my psychological life but is psychologically molecular. It is composed of all, or some, of the 'more basic' elements we have on the list [of necessary conditions] we gave" (2017: 270).

Why resist a psychologically molecular account? O'Brien presses two worries. One centrally involves deviant causation. Leave it aside. The other is a worry regarding circularity. Here is O'Brien:

> You might think it is true that you need to want, know how to, intend, will, when you act. But what is it you need to want, know how to, intend or will to do? The answer, in our case, is 'to raise my arm'. But to raise one's arm is the action we are trying to understand, so to know what all those other conditions are, we need to know what an action is. (2017: 270)

See where this is going. It is not that the agent cannot somehow intend to do something without having a concept of action, though O'Brien's language suggests such an interpretation. The claim is that no account of action is possible that could factor away action into components—action-types must appear as the content of the psychological states driving action-tokens.

This is unconvincing. Above I offered an account of action in terms of controlled behavior driven by motivational (plan-)states. This account need not depend on the notion of action-types in any circular way. For the content of a motivational state could well be the procurement of some desired object, or the production of some stable pattern of behavior, independently of any

notion of an action-type. Once behavior has stabilized—once an agent can produce the pattern with reliability and flexibility in a given set of circumstances—it becomes useful to talk in terms of action-types and tokens. But action itself is not needed to understand the content of every possible plan-state.

We need not take action to be prime. It is composed of elements—in particular, control and a plan.

5.8 Conclusion

We have reached the end of what I think of as this book's first part. Basic elements of agency have been elucidated. We have accounts of control, nondeviance, and intentional action. In the book's second part I begin by continuing to discuss some basic building blocks of agency. But my aim is different. I use this discussion to begin to turn focus away from what is basic, and towards ways that instances of agency can be in excellent form.