

4

Varietals of Control's Exercise

4.1 Control's Exercise

We are now in position to clean up the formulation of control's exercise offered in chapter 2. Recall that I had to rely on an unanalyzed notion of non-deviance to account for the exercise of control. I no longer have to do so. Indeed, it turns out that in exercising the degree of control that she does, the agent engages in non-deviant causation:

EC. An agent J exercises control in service of (aspects or parts of) a plan-state P to degree D in some token circumstance T, where T is a member of a comprehensive set of circumstances C, if and only if (a) J's behavior in T approximates the representational content of (aspects or parts of) P to degree D, (b) J's behavior in T is within a sufficiently normal range for J, relative to J's behavior across C, (c) J's behavior in T is produced by way of those causal pathways that, when taken, lead, with sufficient frequency, to J's reaching D across C, and (d) P is among the causal influences of J's behavior.

4.1.1 A Difference with Bishop

Notice the difference between this claim and John Bishop's understanding of control:

There is exercise of control if and only if the causal link from basic intention to matching behaviour is sensitive, in the sense that over a sufficiently wide range of differences, had the agent's intention differed in content, the resulting behaviour would have differed correspondingly.

(Bishop 1989: 150)

This counterfactual understanding of control's exercise does seem to build in an intuitive understanding of control's possession, by reference to ways

the agent would behave when intending different things. My own account of exercise indexes the exercise to a finer set of circumstances, involving the same plan-state. I would agree with a claim that the agent does not possess much control over her behavior if she could only execute her plan as presently constructed—tweaks to the plan are par for the course.

But control over behavior and control as indexed to a plan-state are subtly different. So, imagine a case in which an agent has a plan, and is well equipped to execute it well across many permutations of the circumstance they find themselves in. But there is a catch—the agent could not execute any of many slight differences in the plan. Certainly this agent is precariously poised in one sense. They should not change their plan. But the exercise of control seems possible for them. Like Dom Toretto reminds his team in *Fast and Furious 6*—stick to the plan.

But doesn't Bishop's account capture something important about control? Certainly. Bishop's account requires the possession of control over multiple plan-states (or if you like, multiple permutations of the same plan). We cannot fully understand that notion without the account of control I have provided. But agents would be in trouble if they did not often possess the kind of sensitivity Bishop targets. This kind of sensitivity is what agents develop as they develop skill at various activities—something I discuss in chapter 7.

4.1.2 Exercising Control over Omissions?

One might wonder how this account of control's exercise squares with omissions—things agents do not (omit) to do. Robinson (2014) has argued that “agents can have at least as much control over their omissions as they can over their actions,” (439) although it seems he would not extend this to unintentional omissions. What about intentional omissions?

In an earlier paper I offered two conditions on (intended) intentional omissions (Shepherd 2014b). (I remained neutral on whether side-effects could qualify as unintended intentional omissions.) First, the omission should match the representational content of a relevant intention. Second, the match should be explained in part by what the intention (or its acquisition) non-deviantly causes (2014b: 20). I explained the causal work of the intention as follows. In cases of intentional omission, the relevant intention (or its acquisition) non-deviantly causes in the agent “a disposition not to A (where non-deviant causation here involves the intention's making changes

to [the agent's] cognitive and motivational systems that are coherently related to the intention)" (23).

I am not going to assess this account here, although I would probably add a control condition. (Something like: the agent should, in the circumstances at hand, possess control over possibilities relevant to the action's non-occurrence.) The question is whether it makes sense to talk of the agent *exercising* control over her intentional omission. In my view, this kind of talk does not make good sense. Those who view omissions as events, or who view absences as part of the causal order, may disagree. I view omissions as absences, and absences as no part of the causal order. Concordantly, I view talk of dispositions to omit to A as elliptical for dispositions the manifestation of which is explanatorily relevant to A's non-occurrence. And, strictly speaking, I do not require these dispositions to manifest—the causal work of the intention is simply to dispose the agent.

So I would rather say that agents *possess* control over events relevant to the intentional omission (this much is consistent with Robinson's above-quoted claim), but that in intentionally omitting to do something, agents need not actually exercise control.

4.2 Voluntary Control

In the remainder of this chapter, I wish to use this account of control to explicate related notions of philosophical significance.

I begin with the notion of voluntary control. Apparently, this notion is central to more than one philosophical debate. At the intersection of action theory and epistemology, philosophers debate whether we have voluntary control over our beliefs (Chuard and Southwood 2009; McHugh 2014; Helton forthcoming). At the intersection of action theory and ethics, philosophers debate whether moral responsibility requires voluntary control (Adams 1985; Smith 2005; Fritz 2018). And the notion makes guest appearances in areas like the philosophy of psychiatry, where theorists have debated whether, for example, addiction undermines voluntary control (Hyman 2007).

The term is often used as though its content is transparent. Sometimes theorists offer a gloss—usually explicating it further in terms of notions like choice or intention. More often, they offer no gloss at all.

Robert Adams is an exception. Here is what Robert Adams writes about voluntary control:

To say that something is (directly) within my voluntary control is to say that I would do it (right away) if and only if I (fully) tried or chose or meant to do so, and hence that if I did it I would do it because I tried or chose or meant to do it, and in that sense voluntarily. (Adams 1985: 8)

Preliminarily: the term “directly” is in parentheses because Adams is working with a distinction between direct and indirect voluntary control, where the former involves only basic action, and the latter non-basic. I discuss direct control more below, so I leave the distinction aside for now.

This account has an intuitive ring to it. Adams is clearly tracking some features of voluntary control. But as it stands, this account does not work.

First, this account allows for compelled voluntary control. An on-the-nose alleyway criminal shows you a gun and directs you to hand over the money. Assume you have voluntary control in Adams’s sense—if you fully try, you’ll hand over the money. This is not, I submit, an instance of voluntary control. The notion of the voluntary is at odds with compulsion. This is easily fixed by adding a non-compulsion condition. I add one to my account.

The second problem is not as easily fixed. It has been recently noticed by Kyle Fritz, who writes that “Not even a professional basketball player can sink a free throw on every attempt, even though she might try, choose, or intend to do so” (2018: 839). That’s true. The best free throw shooters in the world miss between 5 and 10 percent of the time, at least in gametime conditions. Given the high degree of control they possess, I find it plausible that they nonetheless exercise voluntary control over their shots, and in good cases, over their makes.

The more general point is that human agents are imperfect. With respect to even the easiest actions—walking, articulating a familiar word, etc.—we sometimes fail. So the best candidates for actions over which we might have voluntary control are not actions over which we have failproof control. If we require failproof control for voluntary control, no one will have it.

If we wish to fix this problem, we need to be clearer about control itself, and how control relates to voluntary control. Voluntary control, I submit, is a sub-type of control.

A question: what is the target of voluntary control? What is voluntary control over? As with control, voluntary control may be over anything that is within the agent’s causal remit (cf. Robinson 2014: 439). The notion of behavior in play is broad. The agent may have control over her movements, or mental events, or actions, or over specific components of plans or actions that extend well out into the broader world.

Notice, incidentally, that this account places no temporal restriction on what the agent may voluntarily control. Actions take time, and some take more time than others. Why deny that an agent lacks voluntary control simply because the action she is engaged in takes more than a millisecond? Suppose a nearly omnipotent angel forms a plan. Due to her near-omnipotence, there is vanishingly little chance that she will fail to carry out the plan perfectly. Suppose, for example, she intends to perform some sophisticated mental action that is difficult for minds like ours to understand—some feat of imagination. Suppose the action takes a day to complete. Why claim she lacks voluntary control over this action? I see no good reason. I do, however, see a reason to deny that human agents typically have voluntary control over plans that take long periods to execute. Our powers are limited. The future is unknown. Plans that extend far into the future contain joints and steps that range over very uncertain circumstances. So it sounds odd to claim that a human agent has voluntary control over, for example, acquiring a PhD, when the agent is still an undergraduate.

Since the targets of voluntary control are many, I refer to them with the variable *X*. Voluntary control's possession over *X* is the possession of control in circumstances that contain two riders.

First, there should be no impediments to the acquisition of *X*-relevant plan-states—plan-states that include bringing about *X* as an end or means of the plan. This addition is due to cases in which agents could exercise control in bringing *X* about if only they could come to possess plan-states (e.g., intentions) to do so, but in which agents cannot do the latter because of a phobia or some other blocker.

Second, agents should not be in circumstances that have them being compelled to bring about *X*. Voluntary control's possession requires the absence of compulsion:¹

Voluntary control's possession

An agent *J* possesses voluntary control over *X* across a set of circumstances *C* if and only if [a] *C* is well selected, [b] *C* contains no impediments to the acquisition of *X*-relevant plan-states, [c] compulsion to acquire or execute *X*-relevant plan-states is absent across *C*, [d] *J* possesses a degree of control with respect to the execution of *X*-relevant plans, such that *J* brings *X* about

¹ What is compulsion? That could take a long time to spell out. So the explication I offer is in this way incomplete. For here I am working with an intuitive, and rough, understanding of compulsion.

a sufficient number of times across instances in which J acquires an X-relevant plan.

If we wish to move from voluntary control's possession to its exercise, we place the agent in an actual circumstance, and we set a comprehensive set of circumstances. That is, we build a causal model of the situation, specifying the circumstance-type, and causal parameters of the agent and environment, against the background of how these parameters vary in the circumstance-type:

Voluntary control's exercise

An agent J exercises voluntary control over X in a token instance T if and only if [a] J possesses voluntary control over X across a comprehensive set of circumstances of which T is a member, [b] J is executing an X-relevant plan in T, [c] J exercises a sufficient degree of control in bringing about X.

A few features of the account deserve special mention.

First, the account is sensitive to the possibility of an agent exercising voluntary control while failing to bring X about. This would occur if the agent fails to exercise a sufficient degree of control, thus failing. In that case we would say the agent exercised voluntary control over a different target—trying to bring X about. The agent could also fail to exercise a sufficient degree of control, and get lucky in bringing X about. That might be a case of deviant causation, and again we would say that the agent exercised voluntary control in trying to bring X about, but not in actually bringing X about.

Second, this account is consistent with the fact that circumstances can be variable, and that this can impact an agent's control. One might have lots of control in some circumstances, and not in others. This may be due to the environment, in different ways. It is more difficult to play basketball outside when it is windy. But difficulty is not the only relevant feature. Some tennis players are much worse on clay than on hardcourt. But clay is not a more difficult surface, just a different one. Control variances may also be due to features of the agent. Lack of sleep, high levels of anxiety, and many other features may influence the agent's control. It is perfectly legitimate to select a set of circumstances that holds fixed some of these features. So the same agent may have a high level of control in one well-selected set of circumstances, and a low level in another.

Third, this account indexes control not only to sets of circumstances, but to plan-states. It makes no comment about how plan-states come about,

beyond the rider that they come about in an uncompelled way. But the acquisition of plan-states could be placed directly under the microscope, by making such an acquisition the potential target of voluntary control. Something like this has been done by philosophers who debate whether decisions—events of intention formation—are ever intentional actions (Mele 2003b; Shepherd 2015a).

4.3 Direct and Indirect Control

Voluntary control is not the same thing as direct control. But the notions can be fruitfully brought into contact.

The most illuminating discussion of direct control is due to Mele (2017a), who notes that “nothing approaching a full account of it exists” (278). Nor does Mele attempt to provide a full account. But he offers significant guidance. First, he remains neutral between two ways of understanding the target of direct control. The target may be an agent’s action, or it may be events or states of affairs that are intrinsic to an agent’s action. The disjunction need not, to my mind, be exclusive—agents may exercise direct control over both.

Second, Mele offers a plausible condition on direct control: “If S exercises direct control over X, then S does not exercise control over X only by exercising control over something else (or, more precisely, something that does not include X)” (280). This is similar to Adams’s idea when discussing direct voluntary control. The notion of directness may be understood at least partially in terms of basic action, where a basic action is an action an agent can perform without performing any other action.

I resist this understanding, however, because in my view control is prior to action. So direct control over behavior is prior to—and would play a role in explaining—basic action. Instead of speaking of basic action, then, I speak only of direct control, understood in terms of Mele’s above condition.

Third, Mele distinguishes direct control from complete control (which is, it seems, the same as what I earlier called failproof control). The distinction is useful. Consider the following case, which is a paraphrase of Mele’s.

Sol is instructed to press only one of two keys on a keyboard (Q or P). He is to decide which one. Each key has a genuinely indeterministic randomizer such that no matter how hard Sol presses, in some cases the key will fail to fully depress. Mele comments: “So Sol never has complete control over whether he fully depresses the key he has selected and never has

complete control over whether he fully presses the Q key or the P key” (283). I agree. I also agree that in spite of lacking complete control, Sol can exercise direct control in pressing a key.

With that understood, we can distinguish between direct and indirect voluntary control:

Direct voluntary control’s exercise

The agent J exercises direct voluntary control over X in a token instance T if and only if [a] J possesses voluntary control over X across a comprehensive set of circumstances of which T is a member, [b] J exercises sufficient voluntary control in bringing X about, and [c] J does not bring X about by exercising voluntary control over something that does not include X.

Indirect voluntary control’s exercise

The agent J exercises indirect voluntary control over X in a token instance T if and only if [a] J possesses voluntary control over X across a comprehensive set of circumstances of which T is a member, [b] J exercises sufficient voluntary control in bringing X about, and [c] J’s brings X about by exercising control over something else, something that does not include X.

One interesting upshot of thinking of things in this way is that the distinction between direct and indirect voluntary control may come to seem less interesting.

Consider an agent who desperately wants to come to believe something. This is usually discussed in terms of whether the agent can form a belief “at will,” where this is plausibly a stand-in for forming a belief by way of a basic action, or an exercise of direct control. But with an account of voluntary control more fully in view, one may reasonably think that the more important issue is simply how much voluntary control the agent has over the item in question. Why worry if the process takes a few steps as opposed to one, if the control is the same?

Of course there is one reason to worry. For human agents, multi-step processes invite more opportunity for failure. Such processes may thereby correlate with less control. But not always. So agents may well have voluntary control over certain items, such as formations of belief, even if they lack direct voluntary control over them. The issue, in the case of belief, is of course partially empirical. It does not seem to me that we have that much control over our beliefs, although we may in some circumstances (Shepherd 2018b). But there is no guarantee, absent empirical details, that

we will have more control over an item simply because we have direct control over it.

4.4 What Is “up to” an Agent

I wish now to extend these thoughts on voluntary control to a further notion, one that haunts the free will debates. This is the notion of what is “up to” an agent. Many find the following claim at least intuitive: An action *A* is free only if it was up to the agent whether she *A*-ed. And yet, as Seth Shabo notes in an illuminating discussion, “the ‘is up to...whether’ locution and its cognates have largely escaped close examination” (2014: 379).

Much discussion of what is up to an agent focuses on the moment of decision (or choice)—the moment of intention formation. This introduces additional complications, for it requires some work to see how events of intention formation could be legitimately considered intentional actions, or exercises of control (see Mele 2003b; Shepherd 2015a). But it is possible to speak of what is up to an agent independently of the moment of decision. It might be up to an agent whether she acts in the way that she does, given some pre-existing intention. Or it might be up to an agent whether she succeeds in *A*-ing. Perhaps, for example, if an agent had exerted more effort, or paid more attention, she would have succeeded.

Here I wish to make a suggestion. What if we understood what is up to an agent in terms of voluntary control? I would not propose an identification of these notions. For an agent may have some low degree of voluntary control over *X*. The notion of what is up to an agent seems a bit stronger. What about this:

Up to an agent (simple reading)

X is up to an agent *J* in some token circumstance *T* if and only if *J* possesses a sufficiently high degree of voluntary control over *X* in a well-selected set of circumstances *C* of which *T* is a member.

We understand a “sufficiently high” degree of voluntary control over *X* as the possession of a degree of control with respect to the execution of *X*-relevant plans, such that *J* brings *X* about a sufficiently high number of times across instances in *C* in which *J* acquires an *X*-relevant plan.

I do not use the term “whether” above, but we could add it in without, it seems to me, doing any violence. So it is up to an agent whether X when the agent is in a circumstance without the compulsion to bring X about, in which there are no impediments to acquisition of an X-relevant plan, and in which the agent possesses a sufficiently high degree of control over the X-relevant aspects of the plan.

This seems at least a plausible explication of what is up to an agent. Thus explicated, however, some philosophers may wonder whether it applies to the usage of “up to” in the free will debates.

Some use the notion of what is up to an agent in a more expansive sense. Shabo (2014), for example, wants to illuminate the kinds of alternative possibilities that are relevant to moral responsibility. He thinks, like many others think, that not every alternative possibility will do. The alternative must be robust. And Shabo understands a robust alternative possibility as one on which it is true to say of the agent that it is up to her whether she realizes it.²

One might place an alternative possibilities reading of “up to” claims against the one I have so far advanced. One might claim, for example, that for any action option A, agent J, and time t, it is up to an agent whether A only if it is up to J whether A or B at t, where B is a second option (either an action, or an omission (perhaps an omission to A), depending on how strong one wants the alternative possibilities to be).

This seems too strong to me, at least in the sense that it does not sound incorrect to my ear to claim that X may be up to an agent even in the absence of robust alternative possibilities. Something like this may be true of Frankfurt’s cases of volitional necessity (Frankfurt 1988), in which an agent’s motivation in some instance is so strong and so clear and so central to how she sees the world, that she can do no other than X-ing at that time.

² Shabo also claims that the “up to... whether” locution introduces an intensional context, and that this places a kind of epistemic requirement upon the agent. Consider Tom, who does not know which ticket is the winner, but who wishes to pick the winner. Shabo writes:

Consider the inference from ‘It’s up to Tom whether or not he picks the ticket on the left’ and ‘The ticket on the left is the winning ticket’ to ‘It’s up to Tom whether or not he picks the winning ticket’. This inference is invalid; it follows only that it’s up to Tom whether or not he picks what is in fact the winning ticket. (2014: 386)

I agree with the point about intensional contexts. This part is covered by my requirement that the agent’s voluntary control be indexed to an aspect of a plan-state. I do not here take on board Shabo’s additional epistemic requirements. For discussion, see Shabo (2014) as well as Kittle (2017).

I could be wrong about this. But for now let us admit two readings of “up to” claims. The simple reading does not require robust alternative possibilities. The robust reading does require them. We can understand both in terms of voluntary control:

Up to an agent (robust reading)

X is up to an agent J in some token circumstance T, and at some time t in T, only if the simple reading applies to X at t, and if in addition, the simple reading applies regarding a second option Y for J at t.

Neither reading says anything about determinism. Both incompatibilists and compatibilists can help themselves. One upshot of the discussion, then, is the ability to distinguish between these readings, both of which may be relevant to free will, depending of course on ancillary considerations.

A second upshot is the ability to discern what is being claimed when we link “up to” claims with free will. Consider the following options.

- [1] An agent J A-ed freely (at time t) only if the simple reading applied to J's A-ing at t.
- [2] An agent J A-ed freely (at time t) only if the robust reading applied to J's A-ing at t.

Both [1] and [2] are to be understood in terms of voluntary control. Both could be strengthened into claims of necessity and sufficiency. The advantage is that we now know what we are assessing when we assess such claims.

I sense that my friends who focus more heavily on free will might find my explication of voluntary control, or of what is “up to” an agent, not quite to the point. I would be interested to hear why. But until then, I offer these explications to them roughly in the spirit in which Salinger offered *Franny and Zooey* (1961) to his editor—that is, the spirit of a one-year-old “urging his luncheon companion to accept a cool lima bean.”

4.5 Conclusion

I have devoted this chapter and chapters 2 and 3 to development of notions that constitute key building blocks of agency. These notions are also critical for an understanding of intentional action. Chapter 5 is about intentional action. I offer an account.