2
Control's Possession

2.1 Introduction

Anything that displays activity, that is a doing of something rather than a mere happening, will involve the exercise of control.[1]

A great sushi chef prepares a plate of food. One sees a series of precise movements. The chef makes accurate, contextually appropriate cuts into a piece of fish. They roll just the right amount of rice into the right shape. They pinch and then spread a finely chosen amount of wasabi. A great sushi chef can do this again and again and again and again.

As she practices, a professional tennis player takes balls and places them exactly where she wants across the court. Over and over she hits a small window on the court. Or, on the run, she lays down a volley with a delicate touch that slows the speed of the ball and that cuts a difficult angle. In the middle of a match this can seem completely improvised. But most successful shots are the product of years of practice.

It is said that if you want to make a certain percentage of three-point shots in basketball—in games, 40 percent is a great number—you should make double that in practice. That means day after day, in a gym alone, you should make eighty out of 100 shots from 24 feet out. And you should definitely shoot more than 100 per day if you want to get numbers that consistent.

The moral applies to any number of action-types—woodworking, painting, playing an instrument, dancing. What practitioners of these activities hone through practice, and what they display, is in part a tremendous amount of control over behavior. They can direct, or guide, their bodies and minds at a fineness of grain in order to consistently produce movements and thoughts that fall within a small space of variability. They do so because they can, and because they want to, and intend to, because these activities

[1] The ideas in the next three chapters are expansions upon, and in many cases refinements of, ideas first floated in Shepherd (2014a). I'm happier with the present versions of those ideas.

The Shape of Agency: Control, Action, Skill, Knowledge. Joshua Shepherd, Oxford University Press (2021).
 DOI: 10.1093/oso/9780198866411.003.0002

set out standards for success that they have learned to achieve again and again and again and again.

I have said that control is necessary for activity—for action, for agency. I have not said that only agents possess control. Engineers and biologists find the language of control useful, and their usage is similar to my own. A system or sub-system with control is a system or sub-system whose behavior can be modeled in terms of goal-states and processes that drive the system or sub-system into states that constitute or cause the satisfaction of these goal-states (cf. Dennett 1984). Such a system may not qualify as an agent.

The trick, for an engineer or biologist, is to understand the joints and levers of the system—to understand how the control is exercised. The trick for the philosopher of mind or agency is to elucidate the philosophically interesting components of controlled behavior and their relations set in the broader context of philosophical reflection on the nature of agents. I want to know what it is for the agent, as opposed to some non-agential system, or some sub-system within the agent—her early visual system, or her circulatory system, or whatever—to exercise control.

2.2 Control's Exercise

When an agent exercises control, they deploy behavior in the service of a certain class of mental states. The class I have in mind may be as narrow as the class of intentions. Or it may be broad. Perhaps desires, urges, various emotional states, states with imperatival content (arguably: pain states), or even certain perceptual states could qualify. That will depend on one's account of the contents and functions of such states. Perhaps packages of states could together qualify. Some think, for example, that an intention is really just a package of a desire and a certain kind of belief (Davis 1984; Sinhababu 2017). Or we could think of control with respect to a package of intentions, sub-intentions, associated beliefs, and so on. I'm neutral on all this.

My requirements: in order to be served by controlled behavior, a mental state or package of states should (a) represent (series of) events, states of affairs, or whatever, as to be done, or eventuated (that is, should set out a goal) (b) play a causal role in the production of (or, at minimum, attempts to produce) the thing to be done (that is, M should move the agent towards the goal) (c) qualify as a state or states of the agent, as opposed to some sub-system of the agent. Notice: (b) requires that the state move the

agent in the right direction, towards the goal. I require that this not be accidental. The state (or package) that moves the agent towards the goal should, then, do so at least in part because the state's content sets out a way to proceed towards the goal.

The third requirement is shifty, invoking as it seems to either a distinction between personal and sub-personal levels (Dennett 1969), or something like a distinction between doxastic and sub-doxastic states (Stich 1978). I rely on an intuitive understanding of states of the agent for now. (More detailed psychological architectures for particular agents would bring into play more detailed criteria for marking the distinction.) On the intuitive understanding, the agent's intentions, beliefs, fears, and so on are at the level of the agent. But states of the agent's early visual system, or states that regulate the agent's updating of long-term memory—states like these do not qualify. So, while the processing in early visual cortex is plausibly controlled processing, it does not qualify as control that the agent exercises. I discuss this issue further at chapter 5.5.3.

I need a term of convenience to refer to the relevant class of mental states. For reasons that will become apparent, I will call them plan-states. When agents deploy behavior in service of plan-states, they aim at success. Success involves a match between behavior and (certain aspects of) the representational content of the plan-state. That is what it is for behavior to be in service of a plan-state. Such behavior is directed towards matching aspects of the content of such a state.

A basketball player intends to make a shot. We can stipulate that the representational content of the intention includes the following plan: "square to the basket, locate the rim, aim just beyond the front of the rim, follow through, make the shot." When the agent is successful, she executes her intention as planned—she squares to the basket, locates the rim, aims just beyond the front of the rim, follows through, and makes the shot.

Talk of a content match between (aspects of) a plan-state and behavior raises the following question: what is the representational content of a plan-state, and what aspects of it are relevant? I will call the relevant aspects means-end aspects. I'm building on Myles Brand's work on the content of intentions, on which the content of an intention is a plan. Here is how Brand introduces the idea:

> An intentional action can be a momentous occasion in one's life, such as marrying, or it can be a mundane occurrence, such as showering in the morning; but in all cases, the agent is following his plan when acting. He

> has before his mind, as it were, a pattern of activity to which he brings his actions into conformity. (Brand 1986: 213)

As Brand notes, it is not immediately obvious what plans are. We sometimes talk of them in ways that suggest that plans are psychological states. But we sometimes talk of them in ways that suggest plans are abstract objects (e.g., "There is a plan for world peace but no one has thought of it and no one will" (218)). So perhaps plan-types are abstract objects and agents, via psychological states, token some of these types. Settling the ontology of plans here is not necessary.

We need some sense of the structure plans take, as well as of the kinds of content plans can embed.

Regarding the structure plans take, Brand notes that very simple plans need be little more than a linearly ordered series of steps. Plausibly a plan could involve only one step: move left, wiggle finger, scream, or whatever. But complex plans might involve conditional structures, sets of sub-goals, embedded contingency plans specifying what results count as second or third best, and so on. Brand (1986: 219) suggests we model plans as ordered triples, like so:

$$\text{Phi} = \langle A, h, g \rangle$$

Here A is a set of action-types (although one could say behavior-types instead, with action-types as a subset of these), h is a function on A that orders its members in terms of dependency relationships, and g specifies which results (events, states of affairs, or whatever) are the actual goals or subgoals embedded in the plan. Brand argues that g is necessary to capture plan structure because two agents could share A and h while having different goals. "You and I might follow the same recipe in baking a cake, yet act on different plans. The goal of your plan might be to produce a finished cake, whereas the goal of my plan might be to test the recipe; nevertheless, we both performed the same types of actions in the same order" (219).

I agree with Brand that a specification of the goal or goals is important. This specification sets a standard for success. I might be aiming to get the taste of butter to mix with the cocoa just right. You might be looking to give something to your dog for its birthday. We might produce very similar cakes, with my effort largely a failure, and yours a smashing success. The difference is in the goals. Perhaps my performance suffered. Or perhaps I

had a bad plan—that is, perhaps the behaviors in A or the dependency relationships specified by h were poorly chosen or poorly constructed. This suggests a distinction between the quality of performance of the behavior-types a plan specifies, and the quality of satisfaction of the goal(s) a plan specifies.[2]

As important as goals are the dependency relationships between behavior-types. A plan for shaving involves the application of shaving cream and the passing of a razor over the skin. It is important that one happen before the other. Other plans could embed contingency structures, back-up strategies, specification of next-best outcomes, and so on. Many behavior-types could be represented by a plan, with some to be performed only on certain branches of a tree, and only in certain orders, owing to the way the plan orders the importance of the goals to be achieved.

The dependency relationships in a plan thus set up a means-end structure as internal to, constitutive of, the nature of plans. Behaviors are indexed to goals as means, and they are weighted against other possible behaviors given various contingencies.

My claim is that controlled behavior is behavior that, with additional constraints added below, approximates means-end aspects of a plan-state. Brand speaks of the agent bringing actions into conformity with intentions. Lilian O'Brien (2012) speaks of a matching between movements and contents of intentions. The idea here is similar. The aspects at issue are the means an agent takes, the ends in view, and the dependency relationships between various means and various ends. One could pull these apart and speak of an agent's fluency at various aspects in isolation:

Agent A performs the behaviors perfectly, but gets them out of order, and fails to achieve the end.

Agent B performs the behaviors poorly, but gets the order right, and fails to achieve the end.

Agent C performs the behaviors perfectly, and gets the order wrong, and achieves the end.

[2] The distinction between behavior-types and goals is useful, and in some cases necessary to capture the structure and content of a plan. But in many cases it is plausible to think that the goal and the behavior-type share a referent. The goal will simply be to perform the behavior-type as specified.

In one sense Agent C got lucky. Usually, achieving ends requires some level of proficiency at behavioral execution, and at following the steps of a plan. In some cases, however, the end is achieved anyway. And of course in other cases, behavioral execution is perfect, as are the steps of the plan, and the agent fails. Perhaps the plan was risky. Perhaps it contained a fatal flaw. My son recently intended to eat a delicious piece of cherry candy by surreptitiously swiping the candy I held and shoving it mouthward. A flawed plan in one respect. For I held a disgusting cola-flavored piece of candy. So my son failed to achieve his aim.

The relevance of these distinctions is that when we speak of behavior conforming to aspects of a plan, we may have one of many aspects in mind. I will tend to gloss over this, speaking of behavior conforming to or approximating a plan. If it is important, however, we could always be more specific, and speak of plan quality, or of behavior conforming to a particular goal, or a particular means.

In general, then, controlled behavior involves a match or an approximation between behavior and aspects of the agent's plan. In particular, it involves a match or approximation between behavior and the ends (or goals) embedded in the plan, or between behavior and the means as indexed to specific ends. So we can speak of control with respect to a specific end, or a specific means, or with respect to the plan taken as a whole. But events not represented as contributing to the furtherance of the plan are not events under the control of the agent.

Regarding the kinds of content a plan can embed: this will depend upon the agent in question. Regarding humans the question is largely empirical. I say largely because there is a limit on what kinds of content could feature in a plan given the structure plans are supposed to take. Consider an iconically structured visual representation of a scene. Following Green and Quilty-Dunn (2017), this is a representation that meets the following principles. First, "Every part of the representation represents some part of the scene represented by the whole representation." Second, "Each part of the representation represents multiple properties at once, so that the representation does not have separate vehicles corresponding to separate properties and individuals" (11). A representation that meets these principles has very little internal structure, and given the way such a representation compresses information, it is difficult to abstract away from its parts. One lesson to draw from this is that some states, such as a simple iconic visual representation of a scene, may not be able to encode any kind of goal, and may have trouble encoding any structured sequence of behavior-types. (An icon, could,

however, serve as the specification of a goal. But a goal alone does not make a plan.) The psychological states that direct behavior need more than this.

How much more is open for some debate. Philosophers tend to talk of intentions as propositional attitudes. If the content of intentions is propositionally structured, then it is well suited for expressing plans. For propositions are systematic and recombinable, easily capable of representing sequences of behavior-types and of embedding goals. But it is arguable that we should not think of intentions as (exclusively) propositional attitudes (Coffman 2018). And anyway there are probably ways of tokening plans and goal-states without recourse to fully propositional structure. Philosophers have argued that we engage in practical reasoning via non-propositional representational states: map-like representations (Camp 2007), or analogue magnitude representations (Beck 2014), or mental imagery (Gauker 2011), or combinations of these (Shepherd 2018a). As I say, that is an empirical question, and is not my chief focus here.

Plausibly, then, plans can take a variety of representational forms (see Jeannerod 2006). The present point is that in order to exercise control over behavior, an agent needs a capacity to represent a plan for behavior, however simple or complex.

So one's representational capacities are one source of restriction on possessing a plan-state. Are there any others? Some philosophers have debated whether one could try, or intend, to do what one believes impossible (Thalberg 1962; Mele 1989; Ludwig 1992). These are not quite my questions here. I am talking about plan-states generally, and intentions are only one kind of plan-state. Intention possession may have additional restrictions that plan-states like urges, or sensory motivational states, or motor representations, do not. I am asking what is required for a system to possess a plan-state.

I am not asking what is required for a system to possess a plan-state with a specific content—a plan-state to A, for example, where A is an action variable. I have not offered an account of action yet. We are at a preliminary stage.

Even so, a worry similar to the one regarding intending the impossible arises here. Is it possible for a system to possess a plan-state, if the system cannot—lacks the capacity to—execute the plan? This kind of worry does not really arise with respect to simpler systems. If the states a simpler system tokens do not have the function of bringing about behavior that resembles the content of the state, there is little reason to consider the state a plan-state. Of course a system can find itself in unfavorable circumstances,

and token a plan-state that normally leads to success. Such a system could have a plan-state that is, in those circumstances or at that time, impossible for it to execute. We need to ask a more general version of the question. Is it possible for a system to possess a plan-state, if the system lacks the capacity to execute the plan in any circumstance in which the system could be placed?

In more complex systems, systems capable of some degree of delusion or self-deception, it becomes possible to envision a case in which the system has a plan for doing something, where the something is a thing the system cannot, in any case, do. Is that really a plan-state?

I suppose philosophers could disagree about this. Here is what I want to say. It is too strong to require that the system have the capacity to perfectly execute the plan. We should allow that a system can token plan-states that systematically aim too high, for example. What we should require is that the system have the capacity to cause (execute) some part of the plan. That is, in order to token plan-states, a system should have some causal potency. It is a minimal requirement, but a requirement nonetheless.

Causal potency can be understood as those causal powers (or dispositions) by which an agent behaves—causes things. To a rough approximation, an agent's exercise of causal potency can be measured in degrees, by indexing the exercise to a specific plan, or to a part of a plan. We can, for example, define approximation-level and perfect-level potency.

Approximation-level Potency. An agent J possesses approximation-level potency with respect to (means-end aspects of) plan-state P in circumstances C to degree D if and only if for J, P in C can play a causal role in the production of behavior that approximates (means-end aspects of) P's content to degree D in C.

Perfect-level potency. An agent J possesses perfect-level potency with respect to (means-end aspects of) plan-state P in circumstances C to degree D if and only if for J, P in C can play a causal role in the production of behavior that perfectly matches (means-end aspects of) P's content to degree D in C.

The possession of these levels of causal potency is not sufficient for the possession of corresponding forms of control. Consider Frankie:

Batter. Frankie stands in the batter's box, trembling. Frankie tends to strike out, and he has never hit a home run before. Part of the problem is his swing: an ugly, ungainly motion that rarely approaches the ball. In batting

practice, Frankie's coach will put a ball on a tee in front of him. Frankie hits the tee more often than the ball. Even so, Frankie recently saw a film that convinced him one simply needs to believe in oneself. Thus convinced, Frankie eyes the pitcher and whispers to himself, "Just believe, Frankie!" He then shuts his eyes and intends the following: "Swing hard, and hit a home run!" Here comes the pitch. With eyes still closed, Frankie swings hard and connects, producing a long fly ball deep to left field that lands just beyond the fence.

In his specific circumstances, Frankie possesses perfect-level causal potency regarding his intention to hit a home run in the given circumstances. Even so, the home run does not constitute an exercise of control (even if the eyes-closed swing of the bat does, to some degree).

What else does Frankie need? It is tempting to say that Frankie, or Frankie's intention, needs to bring about the home run in the right way. Frankie's swing, which by stipulation was an ugly, ungainly thing, is analogous to a case Al Mele (1992) introduced regarding a philosopher. This philosopher wanted to distract someone, and so intended to knock over a glass. But this intention upset him such that his hand began to shake uncontrollably, thereby knocking the glass over. The philosopher seems to have even less control than Frankie—in both cases the result accorded with the intention, but deviantly.

Consider the following as an account of control's exercise:

EC*. An agent J exercises control in service of a plan-state P to degree D if and only if J's non-deviantly caused behavior approximates (means-end aspects of) the representational content of P to degree D.

There is something right about EC*. First, it rules out cases like Batter as exercises of (high degrees of) control. Second, it is a very plausible idea that the degree of control an agent exercises has to do with the degree of approximation between behavior and plan content. An intention sometimes causes behavior that fails to perfectly follow the plan, and thus fails to perfectly match the content of the intention. Becky intends to make a shot that is all net—that goes in without hitting the rim or backboard. But the ball hits the front of the rim, then the backboard, and drops in. Clearly Becky exercised a degree of control—the shot was very close to all net, so close that it went in. But her behavior failed to perfectly match her intention. (If Becky bet

money on making it all net, this failure will be important.) Assuming that the plan is exactly the same, it seems Becky exercises less control regarding her intention if the shot is an inch shorter, hits the front rim and does not drop in, and even less if she shoots an airball. Third, EC* seems to capture a core truth about control's exercise: the exercise of control essentially includes an agent's bringing behavior to match the content of a relevant plan.

But EC*'s appeal to non-deviant causation is problematic. If there is no non-circular account of non-deviant causation in the offing, then we will rightly suspect that the account on offer is superficial. In effect, EC* will tell us that the exercise of control is essentially a matter of an agent's bringing behavior to match the representational content of a relevant intention in a controlled way.

I think there is a solution to this problem. It stems from reflection on control's possession.

2.3 Control's Possession

The agent exercises control when she behaves in a certain way, driven and guided by a plan and her commitment to it. In order to exercise control, agents must have control.

When somebody does something that seems lucky, we wonder if they could do it again. If they can do it again and again and again, we no longer believe it lucky. We think they have some control over what's going on. Agents that possess control are agents that can repeatedly execute a plan for behavior.

It's one thing to repeatedly execute a plan in very similar circumstances. But the world is capricious. We might want to see if the agent is poised to handle extenuating circumstances as she brings behavior in line with aspects of a plan. If so, the agent possesses flexible repeatability.

In general, an agent in possession of control with respect to some plan-state is an agent poised to repeatedly execute that plan, even in the face of extenuating circumstances.

To illustrate: hold fixed Frankie's intention and suppose a number of things. Maybe the ball comes in 1 mph faster or slower, or an inch higher or lower, or Frankie's muscles are slightly more fatigued, or Frankie produces a slightly different arc of swing. We can vary Frankie's circumstances any way we like and ask: across this set of circumstances, how frequently does Frankie evince the potency he evinced when he hit the home run? The

answer to this question will give us a measure of the control Frankie possesses regarding his intention.

In order to make sense of flexibility and repeatability, we have to specify a certain set of circumstances. This is not necessarily to say that the possession of control is composed (even in part) of extrinsic properties. In discussing her view of causal powers, Rae Langton distinguishes between extrinsic properties and relational properties, as follows: "whether a property is extrinsic or intrinsic is primarily a metaphysical matter . . . whether a property is relational or non-relational is primarily a conceptual matter: it is relational just in case it can be represented only by a relational concept" (2006: 173). As Langton notes, it is natural to view causal powers as both intrinsic and relational: intrinsic because such powers are "compatible with loneliness" and relational because "we need to talk about other things when describing it" (173). This view is available regarding the control an agent possesses.

Many agents are plastic—we lose limbs, muscle tissue, brain cells. Our control is therefore plastic across circumstances. We learn novel ways of performing tasks, and become adept with various tools. Andy Clark claims that our brains are "open-ended opportunistic controllers"—our brains "compute, pretty much on a moment-to-moment basis, what problem-solving resources are readily available and recruit them into temporary problem-solving wholes" (2007: 101). I think he's right. It follows that circumstances impact the amount of control we possess regarding our plans. So the specification of a set of circumstances requires care.

We get viable and interesting measures of control only when the set of circumstances is well selected. A set of circumstances is well selected when we follow principles for set selection that roughly mirror principles for building an accurate causal model of the agent as embedded in a broader causal system that comprises the kinds of circumstances in which we are interested.

So, for example, the set should be sufficiently large. Think of a set of circumstances with only two members: the case in which Frankie hits a home run, and a case in which he misses the ball. This set is not informative: we need a large number of cases before we get any useful information regarding just how lucky Frankie's home run was. A set is sufficiently large when adding members does not substantively impact the resulting measure of control.

Further, the circumstance selector should accurately specify the parameters that are fixed, and the parameters that vary. In some cases the selector

should specify the degree to which the parameters are permitted to vary. The selector might go beyond this, depending upon the purposes of the exercise. One might attach a probability distribution to the operation of certain parameters, for example. How all of this will go will depend to some extent on the selector's theoretical interests, and to some extent on the control-relevant features of the case—the agent's plan, the agent's constitution.[3]

There are many ways to build a theoretically fruitful set of circumstances. This follows plainly from the facts that there are many agent-types and many circumstance-types, and so the ways that these can come together in contexts of plan execution will be several. At an abstract level, given some agent-type (e.g., an adult human agent) as a target, some strategies for building sets of circumstances will be more theoretically fruitful than others. I elucidate this via discussion below.

As we have seen when discussing causal potency, we can specify a level of content approximation regarding a plan-state, or some aspect of a plan-state. Doing this gives us an in principle way to measure an agent's degree of repeatability with respect to any given level of content approximation. An agent J's degree of repeatability DR with respect to some level of content approximation L in a sufficiently large set of circumstances C is given by J's success-rate at reaching L across C, where successes are exercises of causal potency to the relevant level of content approximation or higher. An agent's degree of repeatability (with respect to some level of content approximation) gives us the degree of control she possesses (with respect to that level of content approximation):

PC. J possesses control to degree DR with respect to some level of content approximation L for a plan-state P in circumstances C if and only if J's success-rate at reaching L across C is DR, C is well selected, and P plays a causal role in J's behavior in every member of C.

Bill is throwing darts. Across a set of 100 circumstances, Bill possesses the intention to hit a bullseye. Suppose, now, that Bill hits the bullseye eleven times: his success-rate at this "perfect" level of content-approximation is

[3] Sets can be more or less comprehensive, in covering the range of factors likely to impact success-rates. Agents can possess control even if the set to which we index the measure is not comprehensive. Consider the agent who succeeds at high levels but only in rare circumstances. But comprehensiveness is important to account for. We can say that control has multiple dimensions of gradability, and one of these is comprehensiveness.

0.11. We might focus on other levels of content-approximation as well—the quality of Bill's release, of his stance, of his direction of attention, of his proximity to his target. It is informative to know that Bill places the dart within 1 inch of the bullseye forty-six times, and within 5 inches of the bullseye eighty-two times. We might even change the set of circumstances—adding in various challenging contingencies—in order to measure Bill's control in various ways (I discuss such contingencies below). With each well-selected set of circumstances, and each level of content approximation, we learn a bit more about Bill's dart throwing capacities.

2.3.1 Control with Respect to a Goal

Depending on our purposes, we may wish to hold fixed various aspects of the agent's plan. To see why we might want to hold fixed a goal (or end) embedded in a plan, consider two agents playing darts: one a recreational player (Torrey) and the other a professional (Bill). Torrey thinks he is good, but is he really? No. Especially not compared to Bill. Both possess intentions with this goal embedded: "Hit a triple-twenty." Knowing only this much, we can predict that Bill and Torrey will differ mentally. As they step up to throw, they will likely possess different plan-states. Due to the time spent honing his skills, Bill's repertoire of plans for executing his intention is larger, more complex and subtle, and so on. The details need not concern us—perhaps they include Bill's stance, the feel of Bill's throwing arm, the grip position of his fingers, the direction of Bill's attention. What is important for our purposes is that the type and availability of plans plausibly contributes to Bill's comparatively higher degree of control across a wide range of circumstances.

We can accommodate the difference between Bill and Torrey by specifying a shared aspect of their different plan-states—something they are both capable of aiming for, such as the goal to hit a triple-twenty. This allows us to capture the way an agent's representational capacities contribute to her control, and it allows a more flexible measure of control:

GC. An agent J possesses control to degree GC with respect to a goal G in circumstances C if and only if J's success-rate at achieving G across C is GC, C is well selected, and a plan-state that embeds G plays a causal role in J's behavior in every member of C.

Note that what we have done in isolating a goal we could do for a wide range of aspects: particular behavior-types, particular sequences of behavior, particular kinds of circumstance, and more. In other work I articulate a notion of weak motivation and argue that weak motivation impairs an agent's possession of control (Shepherd 2017a). One might just as well consider how agents fare in conditions of strong motivation, or in conditions of normal motivation. Or one could isolate other factors. Circumstances that involve distracted attention are interesting for some purposes. So too are circumstances that involve heightened or diminished perceptual acuity. Sport psychologists are very good at studying behavioral profiles across interesting variations of just these sorts.

Suppose the amateur is pretty good at his pub, with his darts, but that other pubs and other darts distract him, significantly lowering his success-rates. The pro, however, has a high degree of repeatability no matter the venue, and with most brands of darts. This type of example illuminates a potentially useful measure of control's possession, related to the wideness of the set of scenarios under consideration. In general if J can exercise control with a certain degree of repeatability across a wider range of scenarios than K, J is plausibly better (at least in one sense)[4] at executing the kind of plan-state in question than K.

2.3.2 Basic Levels of Control and Multiple Sets of Circumstances

As I have explicated things so far, an agent's possession of (some degree of) control is relative to a plan-state or some of its means-end aspects, as well as to a well-selected set of circumstances. There are a number of theoretically fruitful ways to select sets of circumstances. What is required is not that any particular set of circumstances be selected, but rather that the selection be fruitful.

This perspective on control leaves room for the following kind of situation. The agent is in some particular circumstance—either hypothetically, or actually—with some particular goal, and we want to know how much control she possesses with respect to that goal. It turns out that the

[4] Agents will sometimes train for, or display talent at, performance in specific circumstances. Some agents may be better than most or all others at some action-type in some restricted range of circumstances.

circumstance she is in contains a range of parameters we might hold fixed. The agent is feeling anxious, and her anxiety levels may increase. The agent is feeling fatigue, and she may begin to feel a strong sense of effort as she acts in one way or the other. The agent's environment contains certain obstacles that might impact her performance levels, and it is unclear which obstacles she may encounter. It will be possible, when assessing the agent's level of control with respect to the goal, to select multiple sets of circumstances. And the agent's level of control with respect, say, to goal achievement—here I'm adding overly precise numbers for effect—may be 0.8 across one set, 0.7 across another, 0.26 across still another.

What are we to say about such cases? One option is to claim that some particular set of circumstances—perhaps the union of all the well-selected sets—is more fundamental than others. On this option, there is a basic level of control an agent possesses. In many cases we require a massive set of circumstances to figure it out.

I resist this option. My rationale is as follows. The point of selecting a set of circumstances is to get (or to fix on) good information regarding how the agent is composed, and how she is disposed to behave, given important factors—stress, fatigue, working memory capacity, availability of tools, etc. Combining well-selected sets may provide useful information in some cases. This will be true, for example, when movement amongst distinct factors produces an interaction effect. But in other cases combining sets of circumstances will degrade information by muddling explanatory factors better kept distinct. So in the case in which we combine all well-selected sets of circumstances, we may end up with something like a single number, but at the expense of any detailed explanation for why the number is what it is. The kinds of systems that engage in controlled behavior are often complex, and there are multiple factors that influence performance, and multiple dimensions along which behavior will differ. It is ultimately unnecessary to average over all the sets of circumstances. We can, instead, view the agent's possession of control as a set of data points, indexed to multiple sets of circumstances. The possession of control looks, in the end, more like a multi-dimensional histogram than a single number; it looks more like a frequency distribution of markers of performance across the range of cases, with as many dimensions as are relevant to capture the behavioral profile.

One implication is that we should not expect to have available, in every case, a clean verdict regarding who possesses more control with respect to some goal or plan. Of course we sometimes speak in this way. We say things

like, "As a dart player, Bill has much more control over his placement than Torrey." Or we say, "Bill has more control over his placement tonight than he did last Tuesday." In my view, such talk should not be taken to refer to anything precise. Often such talk builds in assumptions about normal circumstances, or prototypical performances. Sometimes such talk involves useful oversimplification. Just as often I suspect such talk is unhelpfully vague. If we wish to be precise about the control an agent possesses (or exercises), I think the analyses offered above are the way forward. Our normal talk of control should be seen as a loose, sometimes helpful, sometimes confused, way of referring to the more precise understanding of control developed here.

2.4 Underlying Metaphysics

I have spelled out what it is for an agent to possess control. The metaphysician may feel as though I have offered a framework for control while leaving the underlying metaphysics untreated. They may think that control cannot simply be a distribution of success-rates, or behavioral profiles. Something must explain the success-rates and the profiles.

It seems to me enough to say that control is, in any given case, a package of causal properties sufficient to produce the specified performance profile across the specified circumstances. But one may wonder whether more should be said (certainly, more can be said). For example: might the package of causal properties chiefly or exclusively involve or be best identified with abilities, or perhaps dispositions? There are lively literatures discussing both the nature of abilities, as well as the nature of dispositions.[5] Perhaps the underlying metaphysics of control are best thought of along lines taken by some or another view. Or perhaps the notions of abilities or dispositions as most metaphysicians speak of them are altogether inappropriate as candidates for further explication of control.

In what follows I first consider the relation between control and ability, and then the relation between control and disposition.

[5] One might propose powers rather than dispositions or abilities. Many use the term powers as synonymous with dispositions. Some think of powers in terms of dispositions. But there may be usages that carve out a special role for powers. I do not consider the matter here.

2.4.1 Ability

I do not think control can be understood in terms of ability—at least not as typically discussed. This is because as typically discussed, abilities involve action. John Maier (2018a), in the *Stanford Encyclopedia* entry on "Abilities," suggests that abilities are special kinds of powers. Powers are possessed only by agents, and are expressed by the modal auxiliary "can." What makes abilities special is that unlike powers, they involve action. One might have the power to understand Catalan. But if one can speak Catalan, this is an ability, because speaking Catalan is an action.

Control over behavior is prior to, and necessary for, action. So the relation between control and ability will be one on which the direction of explanation is from control to ability. I do not offer a full version of such an explanation here. The main reason is that there are several candidate accounts of ability in the literature, and this is not the place to decide which ones are better. But it may be useful, for purposes of illustration, to suggest how such an account could go.

Consider a new dispositionalist's analysis of ability:[6]

> S has the ability at time t to do X iff, for some intrinsic property or set of properties B that S has at t, for some time t' after t, if S chose (decided, intended, or tried) at t to do X, and S were to retain B until t', S's choosing (deciding, intending, or trying) to do X and S's having of B would jointly be an S-complete cause of S's doing X. (Vihvelin 2004: 438)

Simplifying, the idea is that an agent possesses an ability to A iff the agent A-s if she tries, decides, or intends to A. This account has been criticized by a number of philosophers (e.g., Clarke 2009, Steward 2012); I am not trying to defend it. Observe however that the account as stated can be seen in terms of control. The account claims that an agent has an ability to A iff the agent possesses and retains certain intrinsic properties—certain dispositions—the manifestation of which constitute a successful execution of an intention (or a trying). As many have pointed out (Austin 1956), that seems far too strong. Agents have abilities to A that fall short of successful performance in every instance. A revised version

[6] See also Fara (2008).

of this account of abilities may do better by seeing abilities as graded in the manner of control. Alternatively, one may think of the possession of an ability as the possession of a sufficient degree of control in some well-selected set of circumstances, or in a sufficiently broad set of circumstances (cf. Manley and Wasserman's (2008) view of dispositions). This may not help with the problem that motivates new dispositionalists—the problem of free will—but it may help with the problem of getting the nature of abilities right.

John Maier offers a view of ability that draws on the metaphysics of modality as well as the semantics of generic statements. Maier wants an account that can explain why an unskilled golfer who luckily sinks a hole-in-one does not have the ability to do so, while a skilled golfer has the ability to sink a putt even though she sometimes misses. On Maier's account, "S has an ability to A just in case A is generally an option for S" (2018b: 425). Option is a technical term, meant to capture an intuitive sense in which some action is actually available for an agent at a time. Generally is a technical term too—"an agent generally has an option" should be understood according to the best semantics for generally. However, Maier admits that we may yet lack an adequate semantics, rendering his account "a promissory note" (421, fn. 7). Maier further suggests that the account may be seen as reducing problems regarding ability to a problem regarding the semantics of the "generally" operator.

It is not my aim here to criticize this as an account of ability. I wish only to illustrate how control may help illuminate it. We can, for example, understand options at least partially in terms of control. I find it plausible that control is at least a necessary condition on the presence of an option: an agent has an option to A in some circumstance T only if she has some sufficient degree of success-level control regarding a plan that includes A-ing as a means or an end across a well-selected set of circumstances C of which T is a member.

I also find it plausible that control is a sufficient condition on the presence of an option: an agent has an option to A in some circumstance T if she has some sufficient degree of success-level control regarding a plan that includes A-ing as a means or an end across a well-selected set of circumstances C of which T is a member.

If this is right, then we need not take options to be metaphysically primitive, as Maier does. We can instead understand the presence of options in

terms of the possession of control at a time. In this way control becomes crucial to an account of ability.

Might control's possession be important as well for an understanding of when an agent generally has an option? I suspect it could be, but this is not crucial to the point I wish to make here, which is that ability is the wrong candidate for further metaphysical elucidation of control.

2.4.2 Dispositions and Explanation

Dispositions are more promising as candidates for the metaphysical basis of control.

There are many accounts of dispositions, however. And there is little consensus regarding which of them are likely to be the truth. Some doubt dispositions can be analyzed, and happily deploy the term unanalyzed—e.g., Stanley and Williamson (2017) do this in a dispositional account of skill. Others deny that dispositions are causally efficacious (Armstrong 1968; Mackie 1977). Given the complexity and sophistication of the literature on dispositions, my goal here cannot be to cover the ground (see Choi and Fara 2018 for an introduction). I wish only to illustrate the suitability of the notion of dispositions in this context.

One view of dispositions has it that dispositions are individuated by a stimulus condition and a manifestation condition. Dispositions are fundamentally dispositions to M (to manifest) when S (when characteristically stimulated). So a vase is fragile if it is disposed to break when struck, or dropped, or…fill in the characteristic stimuli. This view can be formalized into the Simple Conditional Analysis: An object is disposed to M when S iff it would M if it were the case that S. Thus formulated, this view of dispositions is subject to numerous counterexamples (Martin 1994; Bird 1998).

Problems with the simple conditional analysis give rise to more sophisticated accounts (see, e.g., Choi 2006, 2008). Two problems for more sophisticated accounts are salient here. First, might an object or system have a disposition even if that disposition does not manifest in every relevant case? And second, how should we specify what it is for some system to be more disposed to M than another?

Manley and Wasserman (2008) note that in spite of problems with conditional analyses, there seems to be some connection between dispositions and circumstances. In some cases, at least, the relevant manifestation conditions seem built in:

> It is hard to believe that there is no interesting connection between conditionals and ordinary dispositional ascriptions. (This would be especially surprising if we accept that some explicit dispositional ascriptions—the highly specific ones—are actually equivalent to conditionals.) (2008: 73)

Manley and Wasserman propose an account that indexes dispositions to C-cases, where these are precise specifications of values for relevant parameters. So, as they say, in the case of fragility, we might wish to specify the height from which an object is dropped, the density of the medium the object is in, and other intuitively important parameters. Once we specify the relevant parameters, we understand the possession of a disposition as follows: "N is disposed to M when C if and only if N would M in some suitable proportion of C-cases" (76).

This reference to a suitable proportion allows them to truthfully ascribe dispositions even in cases of failure to manifest. But what of the gradability of dispositions? Here Manley and Wasserman recognize trouble. The trouble is due to their liberality with respect to the size of the sets delineated in their analysis. Here is what they say:

> Because the C-cases in our domain need not be restricted in any way, absent stimulus conditions are not a problem for [our analysis]. We can simply allow that N is loquacious just in case N would talk in a suitable proportion of situations—any situations at all. (2008: 77)

They wish to understand dispositions in terms of a suitable proportion of manifestation across extremely large sets of circumstances. Now, I wouldn't deny the usefulness of the set of all situations for some purposes. But in many cases, we have to do better. My framework for control counsels us to pay attention to the behavioral system in question. Depending on the system, and the behavior, some circumstances are relevant, and some are not. It is no relevant part of a decision-making system's control profile whether it can decide what to do if transported in time past the existence of the planet, or the death of the last black hole. So, instead of offering an account of dispositions in terms of proportion of manifestation across all cases, I think it better to index manifestations to more explicitly regimented sets of cases.

This allows one to say that in some circumstances, a system S is more disposed to M than another system Y, even if Y is more disposed to M in different circumstances, and even if there is no fact of the matter regarding

whether S is more disposed to M than Y across all cases. Consider this case, due to Manley and Wasserman:

> Take two objects. Suppose that humans, with our paltry strength, are incapable of breaking the first, but that it would take little or no effort for a giant to break it. The second object, however, is strange—it will break under very light pressure, but it is impervious to the destructive efforts of giants who are incapable of applying light pressure. Now, it may well be that, given a suitable notion of proportion, these two objects break in an equal proportion of cases involving the application of stress. The trouble is that we are likely to call the second object 'fragile' but not the first, while the giants will have opposite inclinations. Clearly, we are not only concerned with the proportion of stress-inducing situations in which an object would break, because some of those situations matter more to us than others. And these situations are different from the ones that matter to giants. (2008: 78–9)

How to deal with this? Manley and Wasserman propose a context-dependent weighting on cases. On this weighting, for example, Object 2 is more fragile when humans are talking about it, and Object 1 is more fragile when Giants are talking about it. It is not clear to me whether Manley and Wasserman would also want to say that there is no illuminating answer regarding which thing is more fragile, full stop. But this seems reasonable to me. So I think some dispositions may be best understood as a disposition-in-C, where C is a well-selected set of circumstances.

But we need not endorse this view in order to elucidate the connection between control and disposition. Even if dispositions cannot be understood as indexed specifically to certain sets of circumstances, control can be understood as constituted by circumscribed dispositions: dispositions-in-C.

This allows us to explain control in terms of something more fundamental than success-rates, while remaining neutral on the nature of dispositions themselves.

It may be useful to illustrate that this lesson applies even if dispositions are very different than Manley and Wasserman think. Barbara Vetter (2014) rejects the view that dispositions are individuated by a stimulus condition and a manifestation condition. Instead, Vetter proposes that dispositions are individuated by manifestation without characteristic stimuli. "A disposition is individuated by its manifestation alone: it is a disposition to M, full-stop" (2014: 134).

Vetter's motivation here is to replace analyses of dispositions in terms of counterfactuals with an analysis of dispositions in terms of potentialities—a modal notion she develops at some length in her (2015). One can see how this might allow an explanation of the possession of a disposition that is compatible with a failure to manifest on occasion. Potentialities may work like this. But how are we to understand what it is for a system to be more disposed to M than another?

What I find congenial in Vetter is that, once she has rejected the intuitive delineation of manifestation conditions that characterizes many counterfactualist discussions of dispositions, she is free to recognize just how context-relative ascriptions of dispositions may be. It looks like, even if a disposition can be understood independently of some set of circumstances or stimulus conditions, a disposition's gradability cannot be understood independently of some set of circumstances or stimulus conditions.

Vetter characterizes the gradability of a disposition in terms of proportions of manifestation across cases. A case involves a world, a time, and an object:

> What should determine the proportion that makes a person count as irascible, or a disease as transmissible, is not the number of worlds in which the person becomes angry at least once or the disease is transmitted at least once, but rather the number of individual instances of anger or transmission: in other words, of cases. (Vetter 2014: 140)

This is in the neighborhood of my own method of selecting sets of circumstances for measuring control. So, too, is Vetter's line of thought for getting the right cases. She emphasizes variety:

> Plausibly, an object which breaks in a greater variety of circumstances is more fragile than one which breaks only in one very precisely circumscribed set of circumstances, and likewise for other dispositions.
>
> (Vetter 2014: 142)

Of course, deciding which types of circumstances qualify as producing greater variety requires one to be more specific about the way one selects sets of circumstances—something I have stressed regarding control. As before, regarding control, the specificity we need will depend on the system in question, and on the behavior in question. To restate: paying attention to these details may give one a way to offer an account of control's

possession in terms of packages of dispositions-in-C—that is, not in terms of dispositions full-stop, whatever they may be, but in terms of how a system is disposed in specific circumstances. Such an understanding can explain control in terms of dispositions, while remaining neutral on the specific nature of dispositions (and on whether dispositions can be reduced to something else).

Readers may have a complaint with this dispositional understanding of control. For notice that this dispositional understanding claims to further illuminate the nature of control, by linking behavioral profiles to something slightly deeper—a system's dispositions. But are these dispositions really deeper? It may sound trivial to say, as I effectively have said, that when a system J has X-amount of control across some set of circumstances this is because J is disposed to behave in a way that confirms the X-measure across the set of circumstances.

When we are considering one set of circumstances alone, I think this complaint is fair. But the way I have articulated the relation between control and dispositions allows one to offer deeper explanations of behavioral profiles by appealing to dispositional commonalities across multiple sets of circumstances. As I have already noted, in many cases what one needs if one wants a more comprehensive understanding of a system's control, or behavioral profile, is to explore that system's behavior across circumstances that hold fixed different important parameters. So we could, for example, explain J's behavioral profile across one set of circumstances C not simply by appeal to J's dispositions-in-C, but by appeal to J's dispositions across multiple sets of circumstances, including C. C will fit into a broader set that more comprehensively illuminates J's behavioral structure, and by linking J's dispositions-in-C to J's dispositions more broadly, J's control-in-C is given an explanation.

2.5 Conclusion

An agent's possession of control is relative to a plan-state or some of its means-end aspects, as well as to a well-selected set of circumstances, which is essentially a well-constructed causal model that embeds the agent into a particular contingency space. Since there are many ways agents become embedded in these spaces, and since different spaces will test different

aspects of the agent—different dispositions and disposition-clusters—there are many theoretically fruitful ways to select sets of circumstances.

In chapter 3 I build on this understanding of control, and apply some of the ideas to the problem of deviant causation. For it seems that one cannot have control without non-deviant causation. More than one writer has made this connection. As Jesús Aguilar has put it:

> What all cases of deviance exhibit is precisely the undermining of the possibility for action by the undermining of agential control. In turn, this explains why any plausible strategy to deal with the problem of deviance boils down to an effort to restore in one way or another agential control.
>
> (Aguilar 2012: 9)

Aguilar speaks of the undermining of action by the undermining of control. I have said very little about action thus far. Just now I am talking about control. In chapter 3 I demonstrate how an account of non-deviant causation falls out of reflection on control.

But I am ultimately headed for action, and I want my account of control to be of use there. Indeed, I will argue that my account of control plays a crucial role in a satisfying account of intentional action. That happens in chapter 5.