7
Skill

7.1 Introduction

Pistol Pete Maravich was a skilled basketball player. You can look up his numbers. One thing the numbers don't show, however, is the unpredictability of his game. His passes, his shots, the way he moved the ball in space—somehow routinely, he violated expectation. Expectations are currency in basketball. If you know your opponent's expectations, you can plan to violate them. Violated expectations buy time, and create advantage.

Maravich's strange skill set was due in part to his father—the coach Press Maravich. From an early age Press had Pete going through unorthodox drills. The young Maravich may have just wanted to please his dad. But his training regimen instilled a unique set of skills in him. Pete's biographer Mark Kriegel explains:

> The gloves and blindfolds were just the beginning. There were so many other drills. Pete learned the fundamentals, of course: dribbling with either hand, chest pass, bounce pass, foul shots, jump shots, and hook shots. But as the basics could become monotonous, Press invented a more elaborate regimen…In all, there were about forty forms and exercises—'Homework Basketball,' as they would come to be known—to cultivate and harvest every bit of Pete's talent. Press and Pete gave them each names, like 'Pretzel,' 'Ricochet,' 'Crab Catch,' 'Flap Jack,' 'Punching Bag.' He would crouch, his arms moving in a figure-8 motion, between and around his legs, so rapidly that the ball looked as if it were suspended beneath his squatting self. (Kriegel 2007: 64)

Like his dad, Maravich was obsessed with excellence, albeit idiosyncratically. These drills were not geared to produce competence, but rather to push the boundaries of competence towards something better. In this regard Maravich is an exemplar of skill. That's because skill is a mode of agentive

The Shape of Agency: Control, Action, Skill, Knowledge. Joshua Shepherd, Oxford University Press (2021).
 DOI: 10.1093/oso/9780198866411.003.0007

excellence. It is a way that agents, qua agents, display excellence. Any account of skill has to accommodate this directionality.

Of course Maravich is only one exemplar of skill. There are many ways to be skilled, and many things at which to become skilled. I discuss (probably too many) examples throughout this chapter. The space is broad. An account of skill has to provide some unity to the diversity.

I used to have a game called *Infinite Stairs* on my phone. The game involves variations on a simple theme. Using an avatar, one climbs as many stairs as possible as quickly as possible. One does this by directing the avatar to climb stairs in one direction with one button, and in another direction with another button. In practice, this occurs as a series of taps on the screen with one's thumbs or other fingers. The game contains stairs going to infinity, but there are only a limited series of stair combinations that fit on the screen. So one can master, relatively quickly, the relevant combinations—left/right/right/left, or right/left/right/left, etc. One other relevant parameter involves the speed with which one can repetitively hit a single button. And a final parameter involves attention—given the distracting nature of some of the stairs or the background coloring or the motion of the avatar, it is necessary to focus attention away from distractors and only to the upcoming stair combinations. So, when playing *Infinite Stairs*, one improves one's capacity to recognize the relevant stair combinations, to perform the relevant button combinations with speed, and to focus attention in the right ways.

It should be uncontroversial that one can become skilled at *Infinite Stairs*. (After all, I did it.) Other, previously practiced abilities, such as familiarity with a tablet's buttons and with the focusing of attention, will be useful. But one has to structure these abilities in the right way. One has to generate novel visuomotor and visuospatial mappings. So there will be a period during which one is a novice, and then one will gradually improve until at some point, one has become skilled. This might not take very long. In my case, after about a month, the game began to get boring—there were small improvements I knew I could make, but they were very small. I was about as good as anyone should reasonably want to be. And I was about as good as I could be, absent, I don't know, performing quick-twitch exercises with my thumbs.

Other skills take much longer to develop, hold an agent's interest for far longer than a month, and cover more complex territory. Consider the skilled chess player. In cognitive science, this type of person has often served as a paradigm for research into skill—or at least "cognitive" skill—with the result that we know a bit more about skill at chess than in some other areas. Of course, the best chess players in the world today are not humans, and are

arguably not agents. They are special-purpose computers. One reason computers have been able to excel at chess is that skill at chess requires mastery of a very particular body of knowledge (see de Groot 1978; Chase and Simon 1973). It turns out, for example, that chess expertise does not correlate very well with expertise in cognitive subtasks more generally. In one series of studies, chess players were better on a measure of multi-step planning, but not at measures of fluid intelligence, nor at tests for visuospatial and verbal working memory (Unterrainer et al. 2006). Similarly, a psychometric test for chess expertise revealed that the best predictor of chess skill was not any general measure of cognitive skill, such as verbal knowledge or recall, but rather performance in a choose-a-move task. In other words, the best predictor of chess expertise is performance at chess (Van der Maas and Wagenmakers 2005). Instead of fluid intelligence, or problem-solving ability more generally, chess skill may involve something more like the recognitional capacities we deploy when we recognize faces or features of faces (see Boggan et al. 2012).

This is not to say that chess players are unintelligent—clearly chess requires abilities to compute, to imagine, to reason, to recognize patterns, to assess one's own assessments, and so on. However, although this skill depends upon various cognitive abilities, the chess expert need not possess these cognitive abilities to a much greater extent than many others in the population.[1] Instead, expertise at chess can be largely put down to practice at chess. In this, chess expertise is similar to expertise at music, or at visual medical diagnosis (Ericsson and Lehmann 1996). The key is not necessarily general intellectual acuity, but intellectual acuity honed in a specific way, for a specific body of knowledge and set of action-types. This is why, as Ericsson and Lehmann (1996) note, "correlation between IQ and performance in a domain decreases over time, and after more than five years of professional experience the correlation is no longer reliable, even after appropriate statistical correction for restrictions of range" (281). Skill at chess often results from the structuring of cognitive ability sets within the normal human range into skill at chess via the acquisition of chess-specific knowledge.

One interesting feature of some activities is that they are variously composed by a wide and diverse set of abilities to perform more constrained activities. Basketball comes to (my) mind. Some of basketball's sub-skills, especially free throw shooting, have been extensively studied by sport

[1] With that said, a recent meta-analysis did find positive correlations between chess rating and measures of visuospatial, numerical, and verbal abilities (Burgoyne et al. 2016).

psychologists. It turns out that some principles important for explaining good performance at free throw shooting are shared with other aiming activities, e.g., putting in golf, archery, dart throwing, and penalty kicks in soccer (Harle and Vickers 2001; Vine et al. 2011). But of course basketball involves much more than free throw shooting. Indeed, it is possible to be truly great at basketball while truly horrible at shooting free throws (see, Shaquille O'Neal). Since basketball is a game involving different roles filled by people with different body-types and different sets of abilities, there are many ways to be skilled at basketball. Certainly almost all of these ways involve certain basic levels of ability to perform rudimentary physical activities—the ability to move with quickness and pace, for example. But more sophisticated abilities are involved as well. One must be able to catch and manipulate and accurately throw a ball of a certain size, to understand the rules and structures of the game and to discern good strategies for play from poor strategies, to recognize an opposing player's intentions and how these impact one's own strategies and intentions, and so on.

The development of high levels of skill, or even of high levels of control with respect to small components of complex activities like basketball, or violin, or oil painting, can take years even for very advanced agents. One must not only develop a range of abilities and master a range of actions. One must become familiar with complex relationships between circumstances, abilities, and success. Some activities require one to compete against highly intelligent opponents who anticipate one's moves and plan counters. One must develop abilities to deploy these behaviors flexibly and appropriately across a range of challenging circumstance-types.

This is all by way of extensive introduction. The extensiveness is due to the fact that, given the variety of ways an agent can become skilled, the proper diet of examples is important. What I'm doing in this chapter is developing an explanation of the nature of skill—of what it is to possess and exercise skill. The account is quite general, but it sheds some light on how humans come to possess and exercise skill. So it can be fruitfully integrated with relevant empirical work. And it places common elements of skill, like knowledge, in the right place.

7.2 Control

If we are building a skill into some agent, where should we start? I suggest control. That's because one clear element of all skills is control. One cannot

possess a skill S without possessing control over some sufficient range of behaviors involved in exercising S.

I developed an account of control in chapters 2 and 3. Perhaps a brief review of the account will benefit readers of those chapters, as well as folks reading this chapter alone.

I distinguish between control's exercise and control's possession. Control's exercise is displayed at a time, as one behaves or acts in a certain way. Control can be exercised to a greater or lesser degree. Consider Fish, who intends to execute the following plan: fake right, wait for the opponent to move right, then head left and around the opponent. Fish has this plan because he believes it is a good way to accomplish a goal of his, which is to shed the opponent. And so it is. Fish fakes right, the opponent begins to move right, and then, as Fish heads left, his feet get tangled and he trips.

Even though he failed to accomplish his goal, Fish has exercised some control over his behavior. As I put it, he exercised control with respect to some aspects of his plan. His fake right went fine. But Fish did not exercise perfect control—he tripped. How to understand the exercise of some degree of control? We focus on certain aspects of the plan, or upon the total plan, which Fish only partly followed. Fish exercises control to the degree that his behavior accords with his plan, or with parts or aspects of it. Or rather, Fish's control has to do with the behavior that non-deviantly conforms to his plan.

Control cannot be exercised unless it is possessed. That's where non-deviance enters in. To see why, suppose things unfold like this: Fish fakes right, his opponent knows the fake is coming and leans left, Fish fail to notice this and heads left anyway, but in so doing Fish's feet get tangled up with his opponent's, and after a bout of stumbling, somehow Fish emerges on the other side of his opponent, who now lies on the ground. Fish has accomplished his goal. But Fish got lucky. He exercised very little control, and we would probably say that he did not shed his opponent intentionally. Perhaps this was just one case. But perhaps the problem runs deeper. Perhaps Fish possesses very little control over his behavior. If so, we would expect Fish to commonly make mistakes like this. For the possession of control is about how one would exercise control across various sets of circumstances.

The possession of control is the possession of a package of causal properties that enable the agent to flexibly and repeatedly bring behavior to match the content of some pertinent plan-state, across some well-selected set of circumstances. The possession of some degree of control is a necessary condition on a plausible account of skill. Given that skill is a mode of agentive

excellence, one might think that the possession of high degrees of control is necessary for skill. The way I understand degrees of control, possession of high degrees of control is possession of dispositions to flexibly and repeatedly bring behavior to (nearly) perfectly match the content of relevant plans, across reasonably large and well-selected sets of circumstances. So, if Fish has high degrees of control at executing post moves, Fish tends to pull off exactly the moves he has planned.

One might now wonder whether high degrees of control are not only necessary, but sufficient for skill.

Say that Fish has a high degree of control at executing the moves he plans. Fish might still contain a fatal flaw. Perhaps his plans are no good. That is, perhaps the perfect execution of the plans he concocts are rarely conducive to the success of his basketball team. They may look nice, or be impressive as one-off athletic feats, or even lead to some success on occasion. But his coaches may in the end decide that they cannot play Fish, opting instead to play teammates whose control, while slightly lower, is nonetheless executed in service of plans that help the team.

Even so, one might think that Fish is skilled at something. He can perfectly execute his plans in a flexible, repeatable way. Certainly something has gone wrong with Fish—he doesn't seem to understand basketball. But something has gone right as well. Fish is very good at executing certain actions within the broader space of the sport.

We might say that what Fish has is skill at various action-types. This notion deserves brief elucidation.

7.3 Skill at Action

Skill essentially involves an agent's being excellent in some way. The skilled agent is skilled—possesses skill, exercises skill—at something. Normal talk is permissive about the boundaries of this something. Philosophers tend to talk of being skilled at an action-type. Indeed, most of the extant accounts of skill bound skill in exactly this way—an account of skill is an account of skill at A-ing, where A-ing is an action-type (e.g., Stanley and Williamson 2017).

If that's all we have to explain, then it's possible we're almost done.

Assume juggling is a single action-type, with whatever sub-actional components (tossing, catching) necessary to fill out the story. What is it to

be skilled at juggling? I suggest that to be skilled at juggling is to possess sufficiently high levels of control with respect to good plans to juggle. This would decompose into the possession of sufficiently high levels of control at sufficiently many of behavioral components of juggling. And this lends itself to a clean account of degrees of skill. One's level of skill at juggling increases along three dimensions. One dimension is completeness—the percentage of behavioral components with respect to which one has some very high degree of control. Another dimension is control—one's success-rate at plans to juggle, or at the components of plans that have juggling as a key goal, where the successes occur in virtue of the control that one possesses (that is, occur non-deviantly). A third dimension is plan quality—the overall success-conduciveness of the agent's plan for action, given her dispositional structure.

When thinking of action-types with relatively simple structure, it can seem as if control is all one needs to account for skill. This is in part because in the case of relatively simple action-types, it is easy to assume that good plans come for free. Perhaps in some cases they really do. Sometimes only minor familiarity with an action-type is needed to bestow the ability to form good plans for performing it.[2] But Fish's case suggests that this cannot work in full generality. It is possible to have high degrees of control with respect to plans, and actions, which are—in the broader context in which they are embedded—counter-productive.

Reflection on skill's variety suggests skills often run wider and deeper than talk of action-types alone can capture. A skilled debater is good at various kinds of reasoning, at listening, at a way of speaking, at synthesizing information, at presenting information. A skilled surgeon possesses a high degree of dexterity of hand and fingers, coupled with a refined understanding of the function of some part of the body, the ways this function may break down, the ways it may be repaired, as well as an ability to apply this understanding to a variety of case-types: to micro-differences in injury and damage and body-type.

The skilled agent is typically skilled at clusters of action-types. These clusters tend to hang together in a structured way. In order to understand skill we have to understand the structure of these clusters—of what I call action domains.

[2] In fact I think cases exist in which an agent can form good plans for A-ing despite having no concept of A-ing, and despite never having A-ed before. I discuss such a case in chapter 5.4.

7.4 Action Domains

The notion of an action domain is familiar from normal talk about skills. We can think of chess, or basketball, or surgery, or parenting, or teaching, or dancing as action domains at which an agent can be more or less skilled. Before saying what it is to be skilled at a domain, I want to illuminate some features of domains, and of how domains come together.

The most basic constituent of an action domain is an ideal of success. The ideal qualifies outcomes of action, including the actions themselves. Sometimes all you need to have an action domain is a particular ideal. Admittedly, this can make an action domain extremely broad. But consider a domain which has, as its only ideal, the complete domination of some range of entities (other agents, for example). As I envision it, this domain has no restrictions regarding how the domination is achieved. This domain is thus something like that of total war: *bellum omnium contra omnes*, war of all against all. In *Blood Meridian*, Cormac McCarthy's character The Judge seems to view all his actions as cantered towards success in this domain. "All other trades are contained in that of war," he says (1985: 246). And he views war as an all-encompassing domain. Consider the following two passages from one of The Judge's soliloquies:

> Men are born for games. Nothing else. Every child knows that play is nobler than work. He knows too that the worth or merit of a game is not inherent in the game itself but rather in the value of that which is put at hazard. Games of chance require a wager to have meaning at all. Games of sport involve the skill and strength of the opponents and the humiliation of defeat and the pride of victory are in themselves sufficient stake because they inhere in the worth of the principals and define them. But trial of chance or trial of worth all games aspire to the condition of war for here that which is wagered swallows up game, player, all. (246)

> It makes no difference what men think of war, said the judge. War endures. As well ask men what they think of stone. War was always here. Before man was, war waited for him. The ultimate trade awaiting its ultimate practitioner. That is the way it was and will be. That way and not some other way. (245)

I doubt The Judge is correct that all games aspire to the condition of war. The present point is that total war is plausibly a domain of action, constructed out of a single ideal of success.

In some cases the ideal is well known. But in many domains, more than one ideal exists. There may be controversy about how best to articulate the ideal, or about which ideals are most important.[3]

Whatever the ideals in play, often they will imply, or will explicitly contain reference to, constraints of various sorts. These constraints can be built out of different elements. Consider the following four elements, common to more familiar action domains: [a] goal(s), [b] an ordering over the goals in terms of centrality-to-success, [c] restrictions on circumstance-types, [d] restrictions on behavior- (and action-)types.

Typically, an action domain will include some restrictions on the circumstance-types that may arise, as well as the behavior-types and action-types an agent may perform within that domain. Many domains are highly restrictive in both ways. Think of stand-up comedy, synchronized swimming, various board games, cricket, as well as our earlier examples of *Infinite Stairs*, chess, and basketball. The presence of some restrictions is the normal case.

Within the space of more familiar action domains, there will often be more than one important goal, as well as an ordering on the importance or centrality of the various goals. This ordering depends upon some axiology for the domain—upon the ideals of success. Some goals are peripheral, and their satisfaction makes only minor contributions to success in the domain. Other goals are critical. In basketball, success is to score the most points. Given the rules of the game, this gives rise to two central sub-goals: scoring, and preventing one's opponent from scoring. Further sub-goals—goals to do with shooting technique, perhaps, or with defensive technique—make fairly direct contributions to the satisfaction of these goals. A goal involving crisp bounce-passes can contribute, but in a less direct way. All else equal, one should practice one's shot more than one's bounce-pass. If you agree, you share some of the implicit understanding I have regarding the way basketball's goals are ordered.

Many domains lack a single master-goal. They may have a central cluster of equally important goals. Or there may be some goals the satisfaction of

[3] Compare the suggestion made by some philosophers of sport that games, in particular, are governed by an ethos (Morgan 2004; Russell 2004; Simon 2000). According to Russell, for example, "games create opportunities for developing certain human excellences by presenting obstacles that must be mastered and overcome in order to achieve the goal set by the game" (Russell 2004: 146). Russell argues that game players should play the game "in such a manner that the excellences embodied in achieving the lusory goal of the game are not undermined but are maintained and fostered" (Russell 2004: 146). I discuss games as a target of skill in (Shepherd forthcoming-b).

which constitutes success, even if they are not the main goal within the domain. In complex domains, much of this will often be hotly contested. What constitutes success at philosophy, or at teaching, or at painting, or at improv comedy? I am not certain, but I have had enough conversations with some of these people to believe that beliefs differ. There may be a pluralism of top-level goals. There may be many routes to success. There may be multiple ways to exhibit skill.[4]

In some complex domains, one way to succeed is to develop sub-skills, or partial skill (as you like): e.g., to specialize at one important thing. In basketball, for example, this may be ball-handling, shooting, or shot-blocking. If one is good enough at one thing, the absence of (high levels of) skill at other elements of basketball can be forgiven. But of course a better way to succeed as a basketball player, if one can manage it, is to develop interesting combinations of sub-skills. These days, the ability to play defense against multiple positions plus the ability to make three-point shots will make one very rich. But look, it is not easy to develop both sub-skills in tandem.

[4] Because action domain construction is often a species of social construction, an additional point about domains deserves mention. The goals and restrictions that constitute a domain often have a kind of internal rationale. This rationale may be (and often is) contested. At times agents not only seek to perfect their skill within some domain. They may seek to tweak or manipulate or change the domain as well. What is it to engage in domain manipulation? I do not have an account, although I find the question interesting. It seems that an agent may manipulate a domain by manipulating constituents of that domain—behaving or arguing in a way that convinces others that new goals or new restrictions are consistent with the domain's rationale. Or an agent could change the way that people view the behavioral space by introducing unexplored behavior-types or action-types.

The acceptance of the changes agents wish to introduce is often a social matter. (Sometimes it happens without prompting by any one agent, of course.) An agent wants to show other agents what is possible regarding a domain, and so seeks to bring about something new. This has often happened amongst skilled artists, when a new technique is invented that changes the scope of what is possible within some genre of art (or that creates a new genre). This happens when musical artists invent a new technique or style of playing an instrument, or when they subtly change the composition of an instrument to make new sounds possible. Or consider how Asher Lev, a painting prodigy and the titular character in Chaim Potok's novel *My Name is Asher Lev*, describes an interaction with Michaelangelo's David:

> The following morning, I returned to the Accademia and stood for more than an hour drawing the David. I drew the head, with the eyes that reflected the decision to enter the arena of power; I drew the huge veined hands that would soon kill; I drew the shouldered sling being lifted in preparation for the delivery of death. The little man with the broken nose had created this sculpture in an act of awesome rebellion against his tradition and his teacher. Other Davids I had seen were small in size and represented David after the battle. This David was a giant and represented the decision to enter the battle. The little Italian had effected a spatial and temporal shift that had changed the course of art. (Potok 1972: 297)

It doesn't seem too farfetched to think that some manipulations of action domains are fruitful, and some less so.

Defending multiple positions calls upon a number of cognitive attributes, such as understanding the typical goals of different players, understanding where each player fits in a team's offensive scheme, and so on. And of course it calls upon physical attributes, requiring, almost as a necessary condition, speed and quickness and a physical frame between 6-foot-5 and 6-foot-9 with a wingspan longer than that. The ability to make three-point-shots involves an ability to reliably throw a ball through a hoop 18 inches in diameter from a distance of roughly 24 feet. What is involved in this reliability is the demand—at the highest level—that one make a good percentage (say, 38 percent) of one's shots across all relevant circumstances, where these include varying levels of pressure, with thousands of people yelling at you, with defenders running at you or sliding under you, while you are moving in one direction or other quite fast just before you shoot, probably breathing heavily and with a massively elevated heart rate, and so on. Here, small differences matter a great deal. Making 38 percent is good, but making 43 percent (something only ten out of 400 players in an insanely competitive professional association obsessed with making that shot can manage) will put one on a different level. JJ Redick keeps getting paid eight-figure salaries in his fourteenth season in the NBA, based largely on his shooting proficiency (this year he is at 46 percent, nearly a career high).

What are the processes that bring the constituents of an action domain together? It seems foolhardy to try to identify some special class of processes. Examining the many types of domains we might implicitly have in mind when discussing skill (human or otherwise), it seems that a pluralism of processes may be at work in any given case. Processes that lead to the development of particular games could come in for further investigation. But games carve out special classes of circumstances and goals, leading to relatively clean domains. Less clean—and in that way probably more interesting—are the domains that many practices and professions set up. The practices of woodworking, painting, sculpture, and playing music set up various domains. So do the professions—legal, medical, educational, etc. Other domains may be broader than these, in the sense that the goals proprietary to them are largely given by biological needs and drives, and the circumstances are largely given by the contexts in which those needs and drives get expressed or find satisfaction. Additionally, it seems that agents may construct novel domains out of existing resources. Many sports and games begin in this way, but domain construction is not limited to sports and games. Some agents may chart new territory in the space of skill by combining atypically combined domains. The skilled politician, for example,

may represent a fusion of the skilled conversationalist, skilled salesperson, skilled self-promoter, skilled debater, and so on.

This diversity of action domain construction is mirrored by a diversity of relationships action domains may bear to each other. Obviously some action domains will be entirely distinct. But many action domains will overlap to some degree. Action domains may share constituents: goals, orderings amongst goals, or various restrictions on circumstances or behaviors. Action domains may also, in spite of not exactly sharing constituents, call for closely analogous ability-types. Consider similarities between rugby, Australian rules football, and American football. Cases of partial overlap present interesting opportunities for the study of human skills. This is because of what skill researchers call transfer of learning (see, e.g., Schmidt and Young 1987). Some abilities, honed in the context of one action domain, may transfer cleanly. Others may only partially transfer, generating a need to fine-tune the ability in new directions—often, a difficult task given the ways human abilities get constructed by practice. Finally, some action domains may be fruitfully thought of as nesting within broader domains. This might be the case for many of the sub-skills I discussed under the heading of skill at basketball. And, although this phenomena is not as useful for thinking about skill, action domains can become entangled or embedded in others.

Recall Pistol Pete Maravich. He was a skilled basketball player. Arguably, he was a better entertainer. You can look up highlights. One thing I see in those highlights is the ambiguous usefulness of many of his moves. If they were designed to win, it's not entirely clear they hit the mark. But if they are designed to amaze, they are perfect.

I'm not the only one to have this thought. Press Maravich was friends and colleagues with UCLA's legendary coach John Wooden. At one point Pete Maravich's biographer has Wooden acknowledge, of a junior high-aged Pete Maravich: "I watched the Globetrotters with Goose Tatum and Marques Haynes. None of them could do more than Pete. Pete Maravich could do more with a basketball than anybody I have ever seen." The reference to the Globetrotters, who play basketball purely for entertainment, is a tell. Wooden was a stickler for playing the game the right way. And he wondered about the aim of Pete's practice. Wooden once asked Press:

> "How many hours does it take to learn all that? Wouldn't he be better off learning proper footwork for defense?"

> "You don't understand," said Press. "He's going to be the first million-dollar pro." (Kriegel 2007: 61)

It turns out, of course, that professional basketball is not all about winning. There are other metrics of success—other domains entangled within the socioeconomics of the thing. Pete's dad Press understood this.

Action domains are fluid. They are not anything like a natural kind, although some domains may prove more central to an agent's survival and flourishing than others. Action domains are often constructed and changed on the fly, based upon an agent's characteristic circumstances or needs or desires. When action domains stabilize, as happens in games and sports, and as happens partially in domains set up by various practices (e.g., various arts and crafts) and professions (e.g., law, philosophy, medicine), this is in part because of a need for coordination of expectations and activities with others.

Given the seeming multifarity of action-domains, and the diverse ways they can be constructed, why think the notion of an action domain interesting or fruitful for reflection on the nature of skill? One reason is that the most sophisticated natural agents we know of—humans—seem to regiment their practice, their habits, and their actions with respect to action domains. This is a natural strategy for any agent that falls short of omnipotence, omniscience, and omnirationality. By constructing an action domain, one is able to carve out a space of goals and circumstances that are manageable, and one is able to begin to map the relationships between the action domain's constituents in a way that facilitates achievement, and ultimately, excellence. A plausible suggestion at this point is that many paradigm human skills are the product of the way that humans carve the space of actions and goals. Humans are often skilled at (or in) some domain, and this is because skills often bear relationships of mutual support to action domains. Both skills and action domains inherit and contribute structure to each other.

A second reason to care about action domains is because there is no clean distinction between action domains and action-types that would warrant the claim that skill at action-types is fundamental. Some action domains could be reconstrued as action-types. Action-types are extraordinarily diverse, and some talk of action-types—baking, dancing, bank-robbing—may just as well be viewed as naming an action domain.

In spite of this last point, the best reason to think of skill in terms of skill at action domains is that this seems the category of appropriate generality

for skill. Skills may be restricted to particular actions. But most of the interesting and paradigmatic forms of skill involve integration of abilities, directed at interestingly clustered action-types and non-actional behavior-types. An account of skill at action alone is thus incomplete.

7.5 Skill at a Domain

We need an account of what it is to be skilled at a domain. Here is mine.

Skill at a domain essentially involves high success-rates at central goals. These high success-rates cannot be due to luck, or accidents. They must, then, be due in part to control. But not only control—at least not only control with respect to the plans an agent may token. Remember Fish's lesson. Good plans do not always come for free. Skill at a domain of much complexity must be due not only to control (that is, the capacity to execute plans), but also to the agent's capacity to (flexibly, reliably) form good—that is, success-conducive—plans.

The plans should be success-conducive for the agent. Some dreamers cook up plans that would be success-conducive for agents who could actually execute them. But dreamers whose plans outrun their own abilities by too much will not enjoy success.

Embedding veridical representations of the world, of nearby possibilities, of one's abilities, is one good strategy for meeting with success. This is, I think, especially true the more complex the domain is in which we are thinking about success. But veridicality is not the only strategy. Plans need not contain entirely veridical representations in order to be success-conducive.

Consider an agent constructed like so. They perceive drops from ledges to be shorter than they in fact are. This could conceivably benefit the development of skill at parkour, if it leads to less trepidation at jumping off of ledges.

Conversely, sometimes it may actually help to aim a bit high. This suggests that plans need not be perfectly executable by an agent to count as success-conducive. Perhaps success-conduciveness is consistent with some degree of self-deception, provided the deception is about the right thing. Human agents tend to do worse under pressure, for example. It may thus benefit an agent if she is able to somehow ignore, or misrepresent, or lie to herself about, the importance of the moment.

The general point is that when rating the goodness of a plan, success-conduciveness trumps veridicality. In many cases the world is such that plans need to embed accurate assumptions about the circumstances. But given divergences between plan veridicality, plan executability, and success, the important feature is the interaction of an agent's plan (or planning style) and the agent's abilities. If a plan (or a planning style) interacts with an agent's abilities to produce a net benefit for success, then the plan (or planning style) is good in that respect.

Here, then, is the picture.

Within a domain, and given some circumstance-type, a space of behavioral routes can be envisioned. These are mappings between possible agents, possible modes of behavior, and outcomes. These outcomes can be scrutinized in terms of the standards for success particular to the domain. The best outcomes will be ones either constitutive of success or else most conducive to success down the way. Depending on the circumstance, and the available outcomes, these will be outcomes that satisfy central goals straightaway, or that bring agents closer to the satisfaction of central goals, where in most domains the satisfaction of multiple goals would be preferable.

Since the diversity of possible agents is very large, the space of behavioral routes is very large. With the domain and constraints on circumstances specified, we can expect patterns to flow through this space. Some routes and some parts of some routes will appear more success-conducive for varieties of agent, while other routes will trend towards failure, or mediocrity. We can imagine toggling settings on the type of agent, and seeing the space change. One set of routes emerges as success-conducive if the agent is very fast on the ground. This set disappears and another emerges when we take away speed, but add flexibility. This will of course be heavily dependent upon the domain, and the circumstance at issue.

Now we focus on a particular agent. Now the massive space of behavioral routes is replaced by a space of available, at least partially executable plans. The notions of availability and executability are generous. These plans are concoctable by the agent—she could token these plans in the circumstance at hand. And the plans are not entirely unrealistic, in that she could execute at least parts of the plans she forms. Viewing this space, again a flow of patterns will begin to emerge. Some plans will seem more success-conducive than others. And we could toggle features of the agent—arousal levels, motivation levels, physical energy stores—and this would shift the success-conduciveness of the available plans.

The space of available, partially executable plans does not characterize the agent's skill. It may characterize something like her (raw) talent. Talent is associated with unrefined potentiality. It is coherent to say that one agent (Zion Williamson) is far more talented, while another (Julius Randle) is far more skilled. And unrefined potentiality sometimes actualizes. A young Edna St. Vincent Millay writes "Renascence," one of her best poems and one of the most memorable poems of the early twentieth century. She then goes to Barnard College. In a biography of Millay, Nancy Milford reports the following:

> William Tenney Brewster's composition classes were a legend at Barnard, where he was a professor of English and provost of the college. Tall and lean, he would sit with his feet coiled about the wastepaper basket, his fingers toying with a rubber band as he read his students' papers in a flat, dry voice. His comments about Millay's work, which were written in a cramped hand on the back page of her themes, were guarded and almost always on mark. He'd given 'Laddie,' about the death of the family dog, a B and said it verged on sentimentality. When she trotted out one of her old St. Nicholas poems, 'Friends,' he wrote 'Browningesque' and gave her a B. And in one of her less inspired themes, when she wrote, 'Why should it be imperative for me to write a theme? System is a fine thing… But even if I were a literary genius (which Heaven forbid) would I be able to—er—give, as it were, whenever System might choose to wiggle her finger at me? Decidedly not,' he marked 'coy' and added to his B, 'Pretty good for the sort; but capable of improvement.' But he continued to encourage her.
>
> (Milford 2001: 100–1)

The young Millay was clearly capable of great heights.[5] But she was an unreliable writer. Her talent was immature. Skill requires more than talent.

So we move from the space of available plans to the agent's performance profile at plan formation and plan execution. We want to see, across a range of circumstances within the domain, how often the agent forms success-conducive plans. And we want to see how success-conducive these plans are.

[5] From her "Renascence": But, sure, the sky is big, I said; | Miles and miles above my head; | So here upon my back I'll lie | And look my fill into the sky. | And so I looked, and, after all, | The sky was not so very tall. | The sky, I said, must somewhere stop, | And—sure enough!—I see the top! | The sky, I thought, is not so grand; | I 'most could touch it with my hand! | And reaching up my hand to try, | I screamed to feel it touch the sky (reprinted in Milford 2001).

The skilled agent reliably winds up with plans that are among the best of those available to her. The more skilled she is, the better the plans she forms. What makes a plan better than another is that, when the agent sets out along the paths it specifies, she is more likely to meet success. The quality of a plan is thus not independent of the agent's possession of control. This is no guarantee that the agent will successfully exercise her control in some particular circumstance. Even the best laid plans, you know.

So:

Skill at some domain D, for an agent J, consists in sufficiently high success-rates for J according to D's standards for success, where J's successes occur in virtue of J's facility at plan construction and J's control over behavior.

In order to have sufficiently high success-rates according to some domain's standards for success, an agent will need to have a facility at plan formation, and at plan execution, leading to satisfaction of at least some central goals in that domain. How much will be required to merit the judgment "J is skilled at D" will depend upon the domain. Skill at *Infinite Stairs* requires much less than skill at basketball, and skill at basketball may come about in a fairly wide range of ways. Skill at teaching is likely to be even more complex, and to permit more diverse manifestations.

Note well, then: skill at complex domains is likely to require elements like knowledge. I discuss knowledge below and in chapter 8. The point here is that skill itself does not require knowledge.

7.6 The Gradability of Skill

Skill is gradable; agents can possess more or less skill at some domain. I propose three principal dimensions along which skill at D may vary.

One dimension is the agent's actual success-rates at goals constitutive of a domain. We might say this is her height at goal satisfaction. All else equal, a higher success-rate at some goal will indicate greater skill. If all else is not exactly equal, there will be a preference for higher success-rates at more central goals.

A second dimension is the agent's success-rates considered across the goals constitutive of a domain. Fix, for example, some success-rate for an

agent. What percentage of the goals in some domain can an agent satisfy at this rate? This is the agent's breadth. All else equal, a greater average success-rate across all the goals will indicate greater skill. In complex domains, all else is rarely equal. So there will be a weighting for central goals.

A third dimension concerns circumstances. All else equal, the greater the range of circumstances along which an agent demonstrates good height and breadth, the greater the agent's skill. We could call this the agent's depth. Again, when all else is not equal, there may be a weighting for more common circumstance-types.

In some domains, there may be a weighting in favor of more difficult circumstance-types. This may enter in via a domain's conception of success—some games, for example, award more points for goal satisfaction in difficult circumstances. In many cases, I believe, the role of difficult circumstances can be understood in terms of success-rates more generally. Difficulty is impressive in part because it suggests that an agent's skill covers circumstances in which many fail, or because it suggests that an agent's skill is especially reliable. Difficult circumstances often offer evidence that is difficult to come by when times are easier.

Variance along any of these three dimensions can indicate greater or lesser skill. But these three dimensions interact. The ideal is an agent who covers all the goals (and, failing that, all the central goals), with extremely high success-rates, across very large sets of circumstances. But cases exist of agents with, e.g., relatively poor height, and relatively good breadth. Such cases may be difficult or impossible to adjudicate, especially regarding domains where the nature of success, or the centrality of goals, is a matter of legitimate dispute.

This way of understanding skill, and the gradability of skill, also illuminates partial skills. An agent's skill may be partial in virtue of excellence along any one dimension coupled with relatively worse performance along the other two. An agent that has an extremely high success-rate at one central goal in a domain is partially skilled. If the goal and the domain are right, this is a good way to make some money—think of long-snappers in American football. An agent that has good height (high success-rates), but limited depth (limited range of circumstances), is again partially skilled. Think of the clay court specialist in tennis. And an agent that has some middling degree of success across a wide range of circumstances, and a wide range of goals, is partially skilled. The utility infielder (Jose Oquendo, a.k.a. The Secret Weapon) in baseball comes to mind.

7.7 How Agents Possess Skill: Skill for Free

I have said very little about how agents come to possess skill. One complaint about the account presses at this point:

El Oso came to boxing later in life. But as soon as he entered the ring, he felt like he was home. El Oso is eight feet tall, weighs 575 pounds, and sports a muscular frame. He's not particularly good at many of the sub-components of boxing. He's not very fleet of foot. His hands are fairly slow. His punches aren't terribly accurate. But there are things El Oso can do. He can take a punch. In fact, he's taken the best punches the best challengers have to give. He shows no signs of damage. And El Oso can throw a punch. They don't always land, but when he's on target, it doesn't much matter if his opponent sees it coming. El Oso will break a rib, or blast right through an opponent's raised hands. The purists don't like it, but El Oso's the champ. He's 57-0, with 57 knockouts. There's no viable challenger in sight.

This story is a dramatization of a claim Stanley and Williamson (2017) make. The claim is that El Oso isn't really skilled at boxing. The reason is something like: El Oso's success is only, or primarily, due to raw physical ability. And raw physical ability is not skill. Here is what Stanley and Williamson say: "Someone's great strength may enable him to win a boxing match despite his lack of skill at boxing" (2017: 717). I agree, but I disagree with the point behind their claim—that factors like strength and speed are, they explicitly claim, "not themselves part of skill." Raw physical abilities make important contributions to skills, and the levels of skill agents possess. They are rarely the whole story regarding skill. But El Oso comes close. What to say?

Distinguish between skill for free and skill in virtue of tuning. Skill for free is still skill. It comes for free because it is attached largely—perhaps in some cases, entirely—to raw abilities. It still requires the possession of some control, and the capacity for constructing success-conducive plans. But in virtue of massive amounts of raw ability, the plans may be relatively easy to construct, and the control required may be undemanding. So it is with El Oso.

I get it. Normal talk of skill tends to reference skill-by-tuning. I am not trying to avoid violence to normal talk of skill at all costs. And it is a cost to neglect the importance of raw ability to skill. In many domains agents need

some combination of raw ability and finely tuned action production capacities. In many domains it is very difficult to come up with an El Oso. Downhill skiing, chess, basketball, gymnastics, philosophy, neurosurgery—in these domains raw ability alone won't cut it. But raw ability certainly helps. And we aren't shy about celebrating excesses of strength and speed and dexterity, or excesses of cognitive control or attentional capacity or foresight in planning. This is because these excesses, which provide the raw material for practice and learning, are themselves important parts of skill. We say of some athletes (Giannis Antetokounmpo) that they are "cheat codes." This is a compliment. To think otherwise is to cleave the agent in two in a way that leaves that agent's successes in exercising her abilities unexplained by the account of skill.

This is not to deny a drive towards fairness, and more watchable competition, in some domains. We manipulate some domains by introducing weight classes, or age restrictions, or whatever. This doesn't undermine the fundamental point.

Nor is this to deny the existence of ancillary reasons for celebrating skill by tuning. Sometimes we track the praiseworthiness of the training an agent undergoes, and the improvements she has made in virtue of her training.

Sometimes these reasons are aesthetic. There is a reason that, when trying to describe high levels of skill, writers reach for terms evoking artistic achievement. So, Paolo Uggetti writes of the guard James Harden, arguably the best guard to play since Jordan: "To call him simply 'methodical' is to do Harden a disservice. He's omni-intentional. Every offensive move seems calculated and artistic, a balletic performance fueled by emotion."[6] When sports writers offer analogies with artistic achievement in more recognizably aesthetic domains (e.g., "balletic"), they attempt to give voice to the fact that high levels of skill in many sports are aesthetic achievements in their own right. This has to do, it seems to me, with the fact that a drive towards excellence coupled with the complexity of many action domains leaves room for, and indeed, seems in some circumstances to call for, aesthetic creativity, the satisfaction of personal goals, and the expression of something like personal style, in exhibiting skill. Here is how Kyrie Irving—widely considered the best ball-handler alive—explains his approach: "A lot of thoughts that you have to put into action…It's just a constant masterpiece that you have to paint. Sometimes it's going to be all scribble and stuff

[6] From Uggetti (2018).

like that. It's okay to get out of the lines."[7] As usual, it's not entirely clear what Irving is saying. But in this case it should be expected. Agents driven to pursue excellence may find themselves navigating an increasingly fine-grained space of reasons for action, and in such circumstances the ability to imprint one's own style on a performance can be an expression of a high level of skill.

Sometimes, of course, in celebrating skill by tuning, we really are tracking a difference in skill. An agent's virtuosity frequently extends her height, breadth, or depth. I would think something like this is true of Lionel Messi, for years now the world's greatest football player. Messi displays virtuosic and complicated combinations of abilities. He is also prone to succeed in situations, and to see avenues for success, that are almost entirely unique to him. We celebrate the way he has tuned his skill for aesthetic reasons, but also because Messi is simply the best.

7.8 Skill and Knowledge

Skill and knowledge are bound up with each other in a variety of ways—more ways than I will chronicle here. My question is fairly specific. What is the place of knowledge-qualifying mind-to-world direction of fit states (paradigmatically: certain beliefs) in the explanation of skill's possession and exercise?[8] Different answers to this question generate two rival accounts of skill.

7.8.1 Pavese's View

According to Carlotta Pavese, skill is intimately related to (cannot be understood independently of understanding) the knowledge of certain propositions. Pavese does not often directly discuss skill—more often, she discusses knowledge how. But two views on skill can be found in her work. Allow me to briefly present both.

[7] Quoted in Chris Forsberg's (2017) story here: http://www.espn.com/nba/story/_/id/21696375/nba-kyrie-irving-again-rises-challenge-boston-celtics-victory.

[8] My question is thus not about the relation of world-to-mind direction of fit states that qualify as knowledge to skilled action. Some have suggested that intentions could be vehicles for knowledge (Campbell 2018; Dickie 2015), even if intentions do not involve beliefs.

In several places Pavese (2016a, 2016b, 2018) links skill and knowledge how as follows:

Claim 1: If one is skilled at Φ-ing, then one must know how to Φ.
Claim 2: Knowing how to perform a task sufficiently well entails that one is skilled at Φ-ing.
Entailment 1: So S knows how to Φ sufficiently well if and only if S is skilled at Φ-ing.

Claims 1 and 2 are presented as fairly intuitive, and do not commit anyone to a specific view of knowledge how. So for all I say here, they may be true. I do not have an account of knowledge how, nor do I want one. But if one wanted to think of knowledge how in terms of dispositions, then Entailment 1 would be unproblematic for my view of skill.

To the above, Pavese adds an intellectualist view of knowledge how, and thereby of skill. But it comes in two strengths. In her (2016a) and (2016b) she speaks of knowledge how "as a matter of" knowing certain true propositions: "According to intellectualism about know how, a subject S's knowing how to Φ, for some task Φ, is a matter of S's knowing a true answer to the question 'How could he himself Φ?' An answer to such a question is of the form 'w is a way he himself could Φ,' for some way w for S to Φ. Accordingly, S's knowing how to Φ is a matter of S's knowing, for some way w to Φ, that w is a way he himself could Φ" (2016b: 650).

Given Entailment 1, this suggests a strong connection between skill and propositional knowledge—skill would just be "a matter of knowing," where the knowledge in question is knowledge of propositional states regarding how to do something. This seems quite strong, but Pavese sometimes speaks of intellectualism about skill as a way of thinking of skill "directly in terms of standing propositional states," (644) and she characterizes the view in one place as the view that "skills are standing propositional states" (647). Call this strong intellectualism about skill.

A weaker view would be that propositional knowledge of the relevant sort is only a necessary condition on skill: S is skilled at Φ-ing only if S knows a true proposition regarding a way to Φ. This is suggested by Pavese (2018)'s endorsement of a belief/knowledge requirement on knowledge how that is only a necessary condition. Call this weak intellectualism about skill.

I am not sure which view Pavese herself currently holds. But I wish to briefly discuss both views.

Pavese frames her account in terms of skill at an action-type. One response to a key fragment of this argument—the fragment behind premise 4—came in chapter 4.4.4. There I rejected Pavese's argument for intellectualism about knowledge how, on which what is required is knowledge as only a necessary condition. I rejected the argument because it takes a belief requirement on intentional action for granted. But I showed that a negative belief requirement is far more plausible. The agent need not believe that w is a way he himself could Φ. The agent must simply fail to believe that the way she acts, w, is a way on which success is unlikely. The rest can be done by the agent's control in the relevant circumstances.

The chief problem with this version of intellectualism is that it needs belief about a way to Φ to be inseparable from the agent's ability to Φ, or control over Φ-ing. But these things are separable, rendering belief otiose in at least some cases.

Conceive of an agent who has never confronted another agent, nor ever considered how to signal anger to another agent, nor considered how to defend their burrow. Then they confront another agent for the first time, as that agent is raiding their burrow. Via specialized perceptual systems, they perceive a threatening agent. This generates a cascade of processes leading to a plan to defend the burrow. (They do not conceptualize the plan as a plan to defend the burrow, of course. But that is what they are doing.) Their brow lowers, their lips thin, their nostrils flair. It turns out this agent is very reliable at defending the burrow. They know how to defend (or, if you would rather, they are skilled at defending) the burrow without having any beliefs about how to do so.

Or, conceive of an agent who has been trained on a stimulus-response set that guarantees high success-rates in circumstances common to some (admittedly, likely very simple) action domain. This agent has no notion that their training was directed towards the domain, nor that it sets them up for success in the domain. Nonetheless, I find it plausible that they possess skill at this domain.

I may be wrong about the plausibility of such cases. Suppose we grant a belief requirement on knowledge how. If so, I'd now wager that the place knowledge has in skill depends upon the agent's control in an interesting way. For what role does the knowledge play? Pavese suggests that it may enter in at the planning stage: perhaps "the choice of appropriate means to ends is itself guided by a standing propositional knowledge state—say, a state of knowing what to do when" (Pavese 2016b: 645). But it is difficult to see how to think of the more determinate content of this state, and of its

characteristic functioning in any particular action domain, without assuming the agent has skill already in place. That is to say, if the knowledge is really going to guide the agent, that is because the agent will be able to deploy the belief in a controlled, non-deviant way. Such an ability does not come for free, simply in virtue of the knowledge's presence in the agent's mind.

I noted something similar in chapter 3.3, footnote 12, when considering Gwen Bradford's (2015) account of competent causation. Bradford argues that an agent competently causes an outcome when that agent causes the outcome while having some "requisite amount" of justified true beliefs about how E is being caused. But, I noted, of course beliefs about the causation of outcomes can deviantly assist the causation of outcomes. One needs a solution to the problem of deviant causation—that is, one needs an account of control—in order to plug beliefs in correctly.[9] The same is true of knowledge. Knowledge can be misused. Knowledge has no magic in the causation of action that intention lacks. It is not enough to posit the presence of a knowledge state. That state must actually guide the agent's action. And it will not do so simply because the knowledge has the right kind of content. The agent must also have control over the use of that state to guide action.[10]

I turn to strong intellectualism about skill. It too is undermined by the above cases, but it merits discussion because, for one reason, in many of the most interesting and complex action domains, beliefs are required for skill. These beliefs will often be even more helpful if they amount to knowledge.

[9] Perhaps doing so would turn the relevant beliefs into knowledge. Dickie (2012) argues for a slightly different version of a skill-explains-knowledge view.

[10] The same is true of packages of knowledge states. I do not commit to any psychological account of how the control is achieved here, but a natural way to go would be to emphasize their structure. So, consider how John Bengson (2017) thinks of the structure of states of understanding that, he argues, contribute to skill. (For Bengson, understanding is a cognitive, epistemically evaluable state distinct from knowledge—but leave that aside.) Bengson argues that practical understanding undergirds manifestations of skill, and that in order to play this role, practical understanding needs several features. Practical understanding of some activity is a conception of the activity in question the content of which is (at least) [a] correct regarding the activity's features, [b] complete in adequately characterizing the activity's central features, [c] internally coalescent in identifying pertinent substantive connections between the activity's central features, [d] externally coalescent in being rationally consistent with alternative conceptions of the activity, and [e] content over which the agent displays mastery. Such a conception, Bengson asserts, is guiding for the agent: "an individual who has practical understanding will be in a state that is action-guiding, poised to underlie and explain the intentional execution of intelligent action" (43). I find Bengson's work here fruitful for further reflection. It seems particularly interesting to think about what sorts of practice, what sorts of mechanisms, and what sorts of capacities of thought might help realize (and to what degree) these properties of internal coalescence and mastery. (In this connection, see Mylopoulos and Pacherie 2017.)

So I think knowledge is rife in skill, even if it is not necessary for skill at an action, or at some action domains. But it is important to see that even in these domains, strong intellectualism cannot be true—skill is not simply a matter of standing propositional states.

Recall the claim that an agent is skilled at A-ing if and only if she knows how to A sufficiently well. If skill is just a matter of intellectualist knowledge how, then we should expect the agent's degree of skill to vary in lock-step with the agent's degree of intellectualist knowledge how.

To assess the viability of this view, we need an account of the gradability of knowledge how. Pavese has done interesting work on this very issue.

Pavese (2017) considers two ways to think of the degrees of what an agent knows how to do. One way, quantitative gradability, involves ascriptions of knowing in part how to A. Pavese offers a picture on which an agent knows in part how to A when that agent knows all of the propositions that are part of the answer to a question regarding how to A. And an agent knows in part how to A when that agent knows some of the propositions that are a part of the answer. One might think of the quantitative gradability of knowledge how as one dimension along which skill at a domain may vary. One agent may know in full how to A for many important action-types within a domain, and only in part how to A for others.

If this is how we think of knowing how to A sufficiently well, however, there are cases that force apart knowledge how and skill. The cases build upon considerations about the structure of action plans, and about the importance of the interaction between an agent's ability and her plans. They involve two agents, J and K. Both know in part how to A, in Pavese's sense of know how. J knows how a bit less—J knows less of the propositions regarding how to A than does K. But J is more controlled, and more successful, at A-ing than is K, in my senses of controlled and successful. How does this happen? This could be due to J's higher control at exercising movements that are important for success at A-ing. Or it could be due to the fact that the propositions J knows, while less than the number K knows, are far more important for successful A-ing across a wide range of circumstances. Or it could be due to the fact that, while J shares the same true beliefs as K, some of J's beliefs have been Gettiered.[11] I would submit that in such a case J is more skilled at A-ing than K, in spite of K knowing better how to A.

[11] It does seem like agents could possess beliefs that are enormously helpful to action execution, but that are not knowledge, due to Gettierization. There is a literature on this: see Poston (2009), Stanley (2011), Carter and Pritchard (2015), Pavese (2018), Carter et al. (2019).

Pavese's second way of thinking of degrees of knowledge how is qualitative gradability—i.e., "Louis Armstrong knew how to play the trumpet better than any of his contemporaries" (Pavese 2017: 369). Pavese observes that a plausible way of reading this claim involves "better than" modifying "knowing how to play the trumpet." As she puts it: "So, playing the trumpet. Who knew how to do it better than anybody? Louis Armstrong did, that's who" (370). Pavese offers a picture on which the way to think of knowing how better than someone else is in terms of the quality of the answers to relevant questions that one knows:

> s knows how to f better than/as well as s` knows how to f' is true (relative to a context c) if and only if there is a practical answer to How to f that s knows (every part of) (relative to c) and that (relative to c) is better than/as good as any practical answer (every part of which is) known by s` (relative to c). (Pavese 2017: 373)

An initial worry is that, if J is more successful at f-ing than K in spite of lacking knowledge regarding one part of an answer for how to f, while K knows every part of the answer, then J is plausibly more skilled at f-ing even though J cannot know how to f better than k. This might be fixed by dropping the requirement that s knows every part of the answer.

Pavese is here understanding the quality of knowledge how in terms of the quality of practical answers. What explains the quality of practical answers? In response to an ancillary objection, she offers this example:

> Suppose Carla and Ale both know several practical answers to the question How to make ravioli but one of the answers known by Carla is better than any of those known by Ale. One way that answer may be better is by being more detailed and precise; or it may be better by being about a better way of making ravioli (a better recipe); a further way her practical answer may be better is by practically presenting a recipe for making ravioli in a better way than any of Ale's answers . . . a practical sense may be better by being more efficient or simpler, just as certain computer programs can be more efficient than others; or it may be better by being more reliable, just like programs can be more or less likely than others to enable the successful execution of the task. By exploiting this further dimension of gradability for programs, my proposal can also make room for the intuition, voiced by Wiggins (2012: 121–2), according to which one may know how to

> perform a task better because, everything else being equal and under appropriate conditions, one tends to be more successful at the task.
>
> (Pavese 2017: 377)

We are given a few suggestions. The first—more detailed answers—is not necessarily associated with more success at A-ing. The second—being a better way of A-ing—seems uninformative. The third—practically representing a better way of A-ing, in virtue of enhanced efficiency or simplicity—does not quite get us to enhanced success in every case. Action in some domains benefits from more winding (less efficient) paths. Others may reward complexity of practical representations. Pavese's fourth suggestion, that a practical representation is better in virtue of being more reliable, sounds like a potential definition of better in terms of success. But although a knowledge-qualifying practical representation may be more reliable in many circumstances, here I reiterate the point that a plan, or a practical representation, may be reliable without being perfectly veridical, or qualifying as knowledge.

Further, the part of an agent's ability that is independent of her practical representation, or her plan, may contribute significantly to her levels of success or reliability, plausibly changing her level of skill without influencing her level of knowledge how.

The strong intellectualist about skill might respond as follows: If the agent's ability is cognitive, then it must just be further knowledge how. And if it is non-cognitive, then it is not a part of her skill. Pavese considers the following objection to her account of qualitative gradability:

> Could not two subjects possess the same amount and quality of propositional knowledge and yet differ in the degree to which they know how to perform a task? . . . If so, one may be better at a task than another because one's ability to perform the task is superior, independently of what propositional knowledge one possesses. "Ability" here means mental or cognitive ability, not simply strength or fitness. (Pavese 2017: 375)

So strength and fitness play no role in quantitative knowledge how. Pavese continues:

> On the general picture outlined thus far, the [mental] ability component cannot vary independently of the knowledge component, for it is knowledge of the relevant practical answer that endows one with the relevant

> ability and corresponding counterfactual success. Thus, on this proposal, it simply cannot be the case that two subjects have the same relevant kind of propositional knowledge about a task—and in particular, knowledge of the same practical answers—and yet differ in their ability to intentionally perform the task (although, of course, they may differ in their nonmental strength or fitness). (Pavese 2017: 376)

In a footnote she adds motor acuity, that is, the changes due to motor-skill learning that enable an agent to execute an action "with more precision and accuracy" (Krakauer et al. 2019: 651), also makes no difference.

Now, these seem false as claims about mental ability. But I do not wish to argue the point. For even if non-mental abilities, and motor acuity, make no difference to one's level of knowledge how, I submit that these plainly make a difference to skill. Indeed, it may very well be because of these features that a particular practical representation contributes to the agent's reliability, or levels of success.

So skill is not just a matter of intellectualist knowledge how. For, in part, the degrees of knowledge how are not the same thing as the degrees of skill. This is consistent with the thought that knowing how sufficiently well—in the ways Pavese's excellent work illuminates—is critical for understanding the structure of many human skills.

7.8.2 Stanley and Williamson's View

Stanley and Williamson (2017) differ from Pavese. But they too would make knowledge prior to skill: "Skill at Φ-ing is a state whose nature is constituted by the knowledge relation" (721). How so?

Stanley and Williamson argue that a skill is "a kind of disposition to know"—that is, "to be skilled at the action type of Φ-ing is to be disposed to form knowledge appropriate for guiding tokens of Φ-ing" (715). It is important that skill is identified with a disposition. This allows Stanley and Williamson to avoid circularity—a skill is not a competence (or any other skill-seeming ability) to acquire knowledge.

Stanley and Williamson do not explain what guidance ultimately comes to. But the notion is important for their account of skilled action, which piggybacks on the account of skill. They draw a distinction between the direct manifestation of a skill, which is knowledge appropriate for guidance, and the indirect manifestation of a skill, which is action guided by acquired

knowledge states: "any skilled action is guided by knowledge that manifests [in the direct sense] possession of skill at that activity" (718). One can discern, then, two different accounts. Skill is a disposition for certain cognitive changes to occur, leading to the acquisition of knowledge. Skilled action is action guided by knowledge the acquisition of which is a manifestation of skill.

One odd feature of this account, of which Stanley and Williamson are aware, is this. One might have thought that the essential manifestation of skill occurs in skilled action. But since they hold that skill and skilled action are separate things, Stanley and Williamson's account of skilled action makes action an inessential manifestation of skill. Stanley and Williamson demote skilled action's role in an understanding of skill in order to emphasize "what is distinctively mental about skill" (721).

One kind of response to this feature of their account is given by Weatherson: "There's something suspicious about a theory of physical skill that divorces it so strongly from the physical" (Weatherson 2017: 382). I think there is something true in that, though skilled action need not be and is not always bodily. We could restate the point like so. There is something suspicious (i.e., false) about a theory of skill at action that divorces it so strongly from the execution of action.

A further problem for Stanley and Williamson concerns the gradability of skill. My points here are similar to those made in response to Pavese, so I will be brief.

Stanley and Williamson mention three ways their account might incorporate gradability. First, one might become disposed to acquire the guidance-apt knowledge more quickly. Second, one might become disposed to acquire more of the relevant facts in a given situation. Third, one might become disposed to acquire qualitatively better information. (They say nothing about what "quality of information" comes to outside of citing Pavese's work on qualitative gradability. Since we have discussed that, I set it aside.)

I think we can agree with Stanley and Williamson that a disposition to acquire knowledge can be graded. I also agree with them that comparisons of skill are sometimes difficult due to the multiple dimensions involved in the relevant assessments. As they say, comparisons of skill often requires "marrying distinct scales" (Stanley and Williamson 2017: 723). However, the account I offer enables a much fuller sense of why this is so. If skills are skills at domains of action, we can see why difficulties in comparisons of skill often emerge. There is often vagueness in the ways action domains get fixed, leading to verbal disputes. So, arguably of course, Jordan is the

greatest player ever, but LeBron may end up with the greatest career; Nadal is better on clay, but Federer is better on grass; Hemingway's use of terse sentence structure makes for thrilling reading, but his female characters are often flat; Einstein was brilliantly insightful, but would be rubbish at running a high-powered modern physics lab; Francis Bacon's portraits are dark, troubling, and great, but David Hockney's almost whimsical portraits may on the whole stay with you for longer; Philosopher A (no names!) is a lovely synthesizer; Philosopher B has the most devastating counter-examples; Philosopher C's ideas are alluring but good grief does C use some imprecise metaphors. Moreover, the account I offer is able to capture the performative element in skill, explaining how two different performances can exemplify different degrees of skill independently of any question about knowledge acquisition during those performances.

Stanley and Williamson's remarks on gradability are hampered by the fact that they separate the knowledge acquired with its role in guiding action. They say that one's disposition to acquire guidance-apt knowledge may improve if one becomes disposed to acquire this knowledge more quickly. But why think quicker equals better? There is no reason, qua disposition to acquire, to prefer speed. Insofar as the knowledge is guidance-apt, one might think it depends on the context of use. Sometimes quicker is better, sometimes slower—it depends on the kind of action at issue, and accordingly on the kind of knowledge and how and when one needs it. The same point can be made regarding their claim that one's skill might improve if one becomes disposed to acquire knowledge of more facts. But more facts do not always mean more control over behavior. Sometimes more facts swamp or distract attention. The point here is that assessments of skill should be linked to some plausible standard. One natural one is the guiding function Stanley and Williamson regard as crucial. But to make that the essence of skill's gradability pulls against their account, on which the guidance of action by knowledge is only an indirect manifestation of skill. As a result, their account seems to imply cases in which an increase in skill undermines the execution of skilled action. The result is that knowledge's prime value in this context—that it tends to guide action better than some representational state that falls short of knowledge, or that is not a knowledge-qualifying state (perhaps, e.g., an intention)—is not well explained.

This is not to deny the importance of cognition, and indeed of knowledge, for most (if not all) of the most interesting skills. One often does need the disposition to acquire knowledge that can—in conjunction with other states such as intentions and perhaps less safe but truth-apt states like

predictions—guide one's behavior. But it seems more plausible to say that acquiring guidance-apt knowledge is something at which one can be more or less skilled. If so, these dispositions are not themselves skill, though they can be structured in ways that constitute an example of skill, and that contribute to skilled action more broadly.

7.9 Conclusion

My overarching concern in this chapter has been to understand skill. In the broader context of this book, I wish that understanding to fall into place as a mode of agentive excellence. Excellence is a kind of perfection of form. So it helps to understand the form of agency. As I have explicated it, this form is that of a system whose behavior, internal and external, is integrated in a way that enables the application of, and the system's meeting of, behavioral standards. Skill can be seen as the possession of structure by a system that enables excellence according to the behavioral standards that action domains set.

I closed the chapter by discussing rival accounts of skill. These accounts overestimate the role of knowledge. But look: knowledge is clearly critical for many human skills. That is not in dispute. In fact, once we see clearly how skill and knowledge are distinct, there is room for a view on which action that intimately involves knowledge—what I call knowledgeable action—forms a distinct mode of agentive excellence. This is the subject of the next chapter.