A Powers Framework for Mental Action

by

Seth Jeremy Goldwasser

Bachelor of Arts, Auburn University, 2016

Master of Arts, University of Pittsburgh, 2023

Submitted to the Graduate Faculty of the

Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Seth Jeremy Goldwasser

It was defended on

June 13, 2024

and approved by

Kate Stanton, Assistant Professor, Philosophy

Wayne Wu, Professor, Philosophy (Carnegie Mellon University)

Dissertation Director: Colin Allen, Distinguished Professor, History and Philosophy of Science (University of California, Santa Barbara)

Dissertation Director: Robert Brandom, Distinguished Professor, Philosophy

A Powers Framework for Mental Action

Seth Jeremy Goldwasser, PhD

University of Pittsburgh, 2024

Mental actions are things we do with our minds. Consider inferring, deliberating, imagining, remembering, calculating, and so on. I introduce a non-reductive alternative to standard causalist accounts of mental action that understands such action in terms of dispositions for performing mental actions. I call this alternative the powers framework. On the powers framework, habitual and skillful mental actions are themselves infused with practical intelligence by being expressions of the agent's rational tendencies and capacities, respectively. The intelligence exemplified in the performance of habitual and skillful mental actions stems from the agent's having shaped the corresponding tendencies and capacities through training and practice. In this way, mental habits and skills are 'second nature' to us. I substantiate the powers framework by giving an account of imagining as a type of skillful mental action. In particular, I argue that imagination is a power to construct representations and select their contents as a means to performing learned behaviors like pretending, engaging with fiction, predicting others' behavior, reasoning about possibility and necessity, reasoning hypothetically or counterfactually about contingent matters of fact, and even imagining for its own sake. I extend the account of imagining to episodic remembering. I argue that such remembering, considered as a mental action, is a kind of imagining by virtue of the agent's constructing a representation and selecting its content as a means of performing the learned behavior of navigating her personal past.

Table of Contents

List of Tables

List of Figures

Preface

This document is something of an ongoing experiment in pursuing some deeply held convictions of mine by way of exploring an interrelated set of characterizations of intentional action, mental representation, mental kinds, and, of course, mental action. These characterizations each strike me as plausible. And, when synthesized in the appropriate way, they together strike me as providing the beginnings of an account of human mindedness. If only I could synthesize them in that way! Hopefully, the attempt I have made here is a good first stab and will not be my last. When I was an undergraduate at Auburn, Eric Marcus started to sew the seeds of discontent with the idea that our minds are just so much grey and white matter. It was also during that period that I learned about action theory. Even then, I found Davidson and the tradition of causalism in the philosophy of action that he originated wanting. It was not until I was in my 2nd year of graduate school at the University of Pittsburgh, taking a course on mental action with Wayne Wu, that I even heard about the notion of a *mental* action. But it immediately seemed to me that understanding mental action might be the key to understanding what I found problematic about the mainstream characterizations of mind and action. It was not until I came to understand Alison Springle's view of mental representation late in the all-but-dissertation phase that I realized that I at once had been going about this whole thing in the most convoluted, backwards way and that I nonetheless had been completely right to think that understanding mental action was key to understanding human mindedness. It was then that I realized that John McDowell had been right when, three years prior, he insisted that all the features of agency in intentional remembering that I attributed to mechanisms of control in my comprehensive examination paper were really features of rationality.

This mostly linear progression in the development of my thought belies many years of vacillation. I was convinced for a long time that all that was needed to deal with interaction problems, interface problems, homunculus fallacies, causal deviance, and so on was a more detailed understanding of the neurocognitive mechanisms that implement cognition. I was tempted by eliminativism and illusionism during my undergraduate years at Winthrop and Auburn. When it became apparent to me that the science is (ideally) in communication with philosophy in determining what the former is doing, how it does what it does, and what it tells us by doing what it does, I was then convinced that clever enough empirically informed philosophers of mind would resolve the aforementioned problems, fallacies, and puzzles with some clever explication or conceptual engineering and a deep and broad understanding of the science. My way of dealing with the tension between a tendency towards reductive naturalism and acknowledging rational normativity was to tell myself

that someone else was going to integrate the scientific and manifest images. It was not until I started presenting at conferences and publishing that I realized that philosophizing about the mind can involve a not insignificant amount of flying by the seat of your pants. And then a whole new world opened up to me: the methodological strictures as well as the metaphysical and meta-philosophical assumptions taken onboard by hardnosed empirically informed philosophers of mind were not set in stone. I was free to try out different ways of approaching the problems I was concerned with; even approaches I had hitherto found unintelligible. And I did just that. I tried to figure out what an account of mental action would look like from within a framework that is roughly in line with the Pittsburgh School. Until I started working on what came to be the final draft of the dissertation, I found this framework fairly alien. It always struck me as making much ado about a certain person-centric conception of our nature that was not psychologistic but nonetheless talked authoritatively about our psychology in a way that resisted empirical inquiry or reductive analysis. Portrayed or approached the wrong way, attempting to understand the framework can be infuriating. Although I now think I better understand it, I have no idea if I succeeded in my aim of providing an account of mental action within the framework. This document is, in effect, my first attempt.

Similarly, the linear progression and the vacillation belie that the trajectory of this project was all over the place. At first, when I had successfully passed my comprehensive examination, I had no idea what I wanted to write about beyond expressing my conviction that there are such things as skillful and habitual mental actions. I was convinced of this in part because much of our mental goings-on do not seem to be things that we deliberate on, yet, are things we can be held accountable for. And, although I did not really have the words for it at the time, this collision of psychological reality and normativity does not sit well with extant, primarily causalist accounts of mental actions in part because almost all such accounts focus entirely on deliberate, self-consciously chosen mental actions. You might call them in a certain sense voluntary. Some accounts have a broader focus than purely voluntary mental action but still model their accounts on unskilled voluntary bodily action. I thought that there was a common misstep here: the kind of control or causal guidance thought to be necessary for responsibility seems to preclude the possibility that the relevant mental goings-on are not things that we deliberate on. Securing that control or causal guidance seemed to require positing an additional state of intention. And this additional state, in order to control or causally guide an already consciously accessible ongoing process, seemed to require that it at least be*come* occurrent and also be consciously accessible. But then the question arises: whence such an additional state? If not a decision, then what? I did not think most accounts had a good answer to this question, if they had an answer at all. So, substantiating the possibility of skillful and

habitual mental action that really did not require this additional state seemed to me to be a great way to point to a widespread but largely unrecognized problem in the literature. I was still of the mind that some form of causalism could accommodate the gap.

But then the trajectory changed as I dug in. I tried a Goldman-style agent causalist approach on which mental actions were type-individuated by their properties, where the agent fixes those properties. Sometimes such fixing goes by way of an occurrent self-conscious intention. But not always, especially not with skillful and habitual mental actions. I found that this approach was a dead-end. I tried a Frankfurt-style guidance-causal approach on which skillful and habitual mental actions only required agential control as mediated by an occurrent intention whenever things were "going off course." I also found that this approach was a dead-end. My first publication assumed the Frankfurtian approach. And at least one article of mine that is currently under review assumed the Goldman-style approach. It finally occurred to me to question the underlying causalist framework when I received comments on an APA colloquium paper from Alison Springle. We talked a few times about that paper in depth and she convinced me that the issues I was running into were a product of holding onto causalism of one sort or another. So, I tried going the anti-causalist route, to go all-in on Anscombe, Thompson, McDowell, Brandom, Boyle, Marcus, Small, and Springle herself. I found it the most satisfying and illuminating approach with no signs of being a dead-end. What once seemed to me to be a fairly opaque and anti-scientific framework for thinking about mind and action now seems like a perfectly cut, ground, and polished lens through which I can more clearly understand much of what I was interested in.

Beyond the linear progression, vacillation, and multiple attempts at attacking the issue from a number of different methodological and theoretical angles, I also would have not created this document without the guidance of several mentors throughout my undergraduate and graduate experience. Eric Marcus was my first mentor. He was incredibly patient with me and encouraged me to be more critical of a variety of standard-stories and mainstream positions in philosophy. We would sit together during his office hours, working through texts on sincerity and testimony. I was incredibly credulous back then. Anything I read seemed right to me upon reading it. I had not developed any ability to find internal inconsistencies nor developed strong enough convictions about the topic to find what was being said an affront to said convictions. Yet, Eric believed in me. He believed that I could take up a particular position and convincingly defend it. And we worked for two years on what would eventually become my writing sample for graduate school, in which I did just that. I do not think I would have had the opportunity to do a PhD without him.

After Eric, there was the committee for my comprehensive examination: Wayne Wu, John McDowell, and Kate Stanton. Although John's role as a mentor for me was brief, I am honored to have had his help in composing my comprehensive examination paper. His encouragement and advice were always invaluable (although I sometimes only realized this in hindsight). Kate has been supporting me since my 2nd year of graduate school, through my comprehensive evaluation, my dissertation prospectus, and the writing of the dissertation itself. Her mentorship and willingness to help me pursue whatever path I wanted, not to mention her great sense of humor and warmth, helped me feel at home at Pitt. I am indebted to her brilliance: she was almost always the first to get me to consider the (sometimes unsavory) implications of a claim that I was itching to defend. Finally, Wayne was an exemplar of the empirically informed philosopher of mind. He deftly guided me through the seminar on mental action and my term paper. I did an independent study with him that built on the term paper and, eventually, for my comprehensive examination, he helped me develop the idea into something publishable. His guidance directly contributed to my winning the 2021 Philosophy of Memory Essay Prize run by the Centre for Philosophy of Memory and, thereby, to my meeting an incredible community of philosophers of memory. Wayne was there for my dissertation prospectus and stayed on as my external for the dissertation itself. Like Eric, Wayne was incredibly patient throughout all five endeavors. Even as I turned away from causalism about mental action, Wayne was always providing excellent advice about how to proceed along empirically informed lines of inquiry and research and helping prepare me for the market. I am grateful to John, Kate, and Wayne for their excellent advice, for their encouragement, for their honesty, and for keeping me honest.

Then, of course, there is my dissertation committee: Colin Allen, Bob Brandom, Kate, and Wayne. Since I have waxed poetic about the latter two, I will focus on former. Colin and Bob both worked tirelessly as my committee co-chairs and as individual mentors. Like Kate and Wayne, Colin has been a positive influence on me and my work since my 2nd year of graduate school, when I took his seminar on the philosophical foundations of cognitive science. Rather than write a term paper on cognitive science, I got stuck on the notion of normal function in philosophy of biology and Colin encouraged me to pursue what would ultimately become two further publications on the notion as it relates to cancer biology. Colin has always pushed me to improve as a philosopher of science. At the same time, he would reaffirm his confidence in me whenever I succumbed to the temptation to express my imposter-syndrome-driven doubts. Both Bob and Colin pushed me not just as a philosopher but also as a writer. Without them, I would be even more obscurantist, inchoate, confused, and un-parsable a writer than I am now! Most importantly, both Colin and Bob allowed me to

pursue research independently while never cutting me slack unnecessarily. I could always think through whatever I found interesting and come to whatever conclusion seemed to me most plausible. But if my arguments and explanations were lacking (as they often were), both would let me know it and both would encourage me to do better. Colin and Bob always know where to push and how much.

Finally, there is Alison Ann "Sonny" Springle. I met Alison my first week at Pitt. And I knew immediately that she would be an important person in my life. At first, she was a mentor. As a graduate student further in the program and with similar interests, she was able to give me advice on which CNBC classes to take or avoid, what to work towards if I chose to professionalize, and which faculty would be most likely to appreciate what I could bring to seminars and in term papers. She likewise encouraged me and other Pitt graduates to submit to conferences, present our work, and network even while we were in our 2nd or 3rd year. She always saw the best in me. And she always encouraged me to be brave, whether it was skipping out on talks at the CNBC annual retreat to go hunt for strawberries in the woods or sticking to my philosophical guns despite pushback. When she was selected to give me comments on a paper concerning automaticity and control in skilled action for the Central APA, I was terrified. And I was right to be: she saw right through my attempt at a little fix. She would not stand for addressing one homunculus-driven problem by positing yet another homunculus. I was now too far in the program and too deep into the process of professionalization to be treated with kid-gloves. Yet, her critique was fundamentally positive. That is, while she found my paper irredeemably flawed, she trusted that I would come to see that it was for myself. She trusted that I would figure it out. And, eventually, I did. That was the beginning of my turn towards the powers framework for thinking about mental action.

Over the next few months, Alison and I talked more and became closer. We realized that we had walked much the same path through graduate school. We both came in thinking some form of empirically informed metaphysics of mind would help us resolve deep confusions we had each been harboring for some time. It then slowly dawned on us through our time in graduate school that we would not be satisfied by what was on offer. Through graduate seminars in the HPS and Philosophy departments at Pitt, CNBC courses, conferences, talks, etc., we came to realize that no satisfying account was forthcoming. We both liked 4E approaches but both found them intolerably anti-representationalist. We also both found standard forms of representationalism confused. And we both found the Pittsburgh School simultaneously more illuminating and less informative than we'd have liked. Alison had unraveled the root issue by the time I was coming around: what was needed was a new account of mental representation based in an intentional-action-first, knowledge-first approach that

articulated both what we share with animals and what is unique to us. By the time I came to adopt the powers framework, I realized that she had already done most of the work that I thought was necessary. Yet, I could bring the strands that I had gathered into conversation with each other. And Alison encouraged me in doing this. We found a deep harmony between our views and approaches. And along the way, we fell in love. I have never been happier. I have never felt more understood by someone. In Alison, I have found a true life-partner: a best friend, a teammate for life; a doting mother to our cat daughters Arya, Bihui, and Kinky as well as to our dog daughter, Freya; and, of course, a loving wife. I could not have done this without you, Sonny. You've given me the courage to try and be myself. And I can never thank you enough for that.

Finally, I want to thank my family. Your constant material and moral support mean the world to me. I love you, Mom, Dad, and Ben.

Introduction

Mental actions are things that we do with our minds. Types of mental action include reasoning, deliberating, imagining, remembering, attending, mental arithmetic, and so on. The principal explanandum in the contemporary philosophical literature on mental action is the agency expressed in the agent's forming, modifying, or maintaining personal-level content-bearing states or processes. For instance, what is it exactly that makes an agent's summing two numbers "in her head," rotating on imagined figure, or listening in on a conversation in the face of distractors things that *she* does? What is it that qualifies these happenings as actions that she performs? In this dissertation, I aim to provide an account of mental action that stands as a viable alternative to standard causalist accounts. The account I provide is both anti-reductionist and anti-causalist. The account derives from a background metaphysics of action that I call "the powers framework." I draw out this framework from work on rational agency primarily by Small (2012, 2014, 2017a,b, 2019, 2020, 2021, 2022) and work on animal agency primarily by Springle (2019, 2021; see also Grush and Springle 2019; Springle and Humphries 2021; Springle and Buccella 2024). Both are, in turn, inspired primarily by the work of Anscombe (1957/2000) and Ryle (1949).[1] I extend the powers framework to address the question of mental action. On my extension of the framework, mental actions are expressions of mental skills and mental habits. And these skills and habits, in turn, are causal powers of the agent to construct and use representations that stand in or go proxy for what they represent. I call these kinds of representations "discursive." Thus, according to the powers framework, mental actions are a matter of an agent's discursive activity, whether covert or overt. That is, on my extension of the powers framework, an agent's expressing a thought is just as much a mental action of hers as is forming that thought. In some cases at least, such expression and formation might well be the overt and covert sides, respectively, of the same coin.

Before giving an overview of the chapters, I want to more firmly fix the topic of discussion. The powers framework and its causalist rivals each constitute an approach to the metaphysical nature of agency in general and mental agency in particular that differs radically from the approach taken by the other. In fact, as will quickly become apparent to those familiar with philosophy of action, these approaches differ enough in their commitments that it can seem that they end up talking past one another. To head this worry off, it will be helpful to have an easily identifiable target in view when contrasting them and when considering the viability of

[1] Although I draw primarily from Small and Springle, there is a healthy ongoing trend of defending this framework in philosophy of action. For instance, see Thompson (2008), McDowell (1996, 2007, 2010, 2011), Marcus (2012, 2018, 2019, 2022a,b), Boyle (2011).

my extension of the powers framework in chapter 2. At the outset, I follow much of the literature on mental action in restricting focus to the mental actions of neurally-intact adult human beings under normal conditions.[2]

The most intuitive gloss on the subject of discussion is that a mental action is the agent's doing something with her mind. Assuming an independent and at least intuitive grasp of agency, we can use this gloss to draw a rough boundary between the class of mental actions and the class of mental happenings that are not actions. Many if not all instances of imagining something seem to be an agent's doing something with her mind and, so, fall within the class of mental actions (see also chapter 3). At least some instances of remembering something and some instances of inferring something likewise seem at first blush to fall within the class (on the former, see chapter 4). All else being equal, choosing to do something and thereby coming to intend to do that thing seems in every case to fall within the class of mental actions (*cf.* Shepherd 2015, 2023). Likewise, some perceptual attention-involving activities like listening to or out for something and looking at or out for something seem by definition to be things that agents do with their minds. Yet, it does not seem like every instance of attending to something is the agent's doing something with her mind, namely, in cases of attentional capture (*cf.* Wu 2011a). Likewise, neither instances of seeing, hearing, smelling, tasting, nor touching seem at first blush to fall within the class of mental actions.

There are, then, a variety of mental happenings that can be intuitively grouped together as falling within the class of mental actions. However, in order to really fix the topic of dispute between the powers framework and its causalist rivals as it concerns mental action, it should be possible to say in a way neutral to both accounts what it is that ties these happenings together as mental actions. Specifically, I here provide a sketch of a neutral account of mental action and give two (at least) *prima facie* reasons for adopting it. I will return to the account in chapter 2 in assessing the powers framework's account of mental action. The neutral account of mental action is as follows:

> **Neutral Account of Mental Action:**
>
> A happening is an agent's mental action just in case:
>
> 1. that happening is the agent's forming, modifying, or maintaining some token of a type of content-bearing, personal-level state or process; and

[2] I have some things to say in passing about the mental actions of younger human beings in the midst of growing "to the age of reason" in chapter 2 (Anscombe 1957/2000: §4).

2. an aim that the agent has in doing what is described in (1) is to form, modify, or maintain a token of that type of content-bearing, personal-level state or process that bears a determinate content specified by that aim.

By "formation" I mean bringing about or constructing a state or initiating a process. An agent can form an intention to do *A* by deciding to do *A*.[3] She can likewise form visual imagery of a dancing banana in a process of visualizing. And she can initiate a bit of visualizing by intending to imagine a dancing banana. By "modification" I mean changing the contents of a state or process. The agent can modify her imagery of a dancing banana during a single instance of visualization by adding or removing details in the process of construction. For instance, she can rotate the banana or give it a cane and top hat. By "maintenance" I mean keeping a state or process present before one's mind or in conscious awareness. An agent can maintain her attention on one conversation in the face of distractors. She can likewise maintain her visual imagery of a dancing banana.

By "personal-level" I mean a state or process that is properly attributable to the agent as agent and that she can in principle access for report or reasoning (*cf.* Drayson 2014). It is the agent who sees when she is in the relevant state, whereas the corresponding activation of neurons in area V1 is not something attributable to her as agent. Moreover, with the possible exception of blindsight, the agent has immediate access to her perceptual state at least in the sense that she is in a position to report and use in reasoning content "gleaned" straight from her perceptual experience. By contrast, she usually cannot report on or reason about the activity of her neurons in V1. She might be able to do this by use of fMRI or a CAT scanner, assuming she knows how to interpret the data they present and has such data readily available as she is being scanned and looking at things. In such a case, the agent would have indirect access to the neuronal activity partially constitutive of her perception. By "content-bearing" I mean a state or process that functions to represent.[4] I take it that mental states are in part defined by their function to represent. Thus, condition (1) already distinguishes mental from non-mental action by restricting the former to specifically happenings that function to represent. That said, the condition leaves open whether the relevant content-bearing personal-level states or processes must exist entirely "in the head" or can extend to and beyond the body. Call condition (1) the "mental condition."

By the agent's having an "aim" I mean that she is in a position to take forming, modifying, or maintaining a state or process of the relevant kind with some determinate content as a reason for doing what is

[3] "Doing *A*" and its cognates are variables for intentional actions that I am borrowing from Small (2022).

[4] To my knowledge, there are no anti-representationalist accounts of specifically mental action.

described in the mental condition. This reason for action can turn out to be all that the agent is aiming at. That is, the agent can, say, imagine something, infer something, remember something, attend to something, and so on for its own sake. Alternatively, the reason can turn out to be instrumental, relating the resulting mental action to the achievement of some further goal as a means. For instance, an agent may intentionally remember her mother's maiden name as a partial means of writing down the appropriate information on some legal document. Or she can intentionally imagine what it would be like to ride a dragon as a means of better inhabiting the fictional world of *A Song of Ice and Fire*. Condition (2) distinguishes mental actions from those mental happenings that are not actions by restricting the former to happenings that end up satisfying the mental condition for specifically practical reasons (for a similar claim made by way of a very different argument, see Hieronymi 2006). Importantly, this restriction is neutral between the powers framework and its causalist rivals primarily in abstaining from the use of causal language. Call condition (2) the "action condition."

Going forward, I will drop from the Neutral Account talk of modification, maintenance, and process. I will likewise replace talk of a state's being content-bearing and personal-level. So, going forward, I will say that, on the Neutral Account, a mental action is the agent's forming some mental state for a practical reason. I take it that the powers framework and its causalist rivals agree with the Neutral Account on what makes something a mental action, especially in this simplified guise.

Before moving on, there are at least two *prima facie* reasons to adopt the Neutral Account for fixing the topic of discussion. The first reason is that the account goes some way to elucidating the intuitive groupings of the relevant mental happenings that were listed some six paragraphs ago as mental actions. The instances of imagining, remembering, perceptual attention-involving activities like listening or looking, paying attention more generally, deciding, and so on that intuitively seem to be among the class of mental actions are happenings that are assessable for satisfying or failing to satisfy criteria of practical rationality. That is, we can assess these happenings for being effective means to the agent's ends, for whether or not they cohere with her other plans and wider actions, and so on. The instances that intuitively seem not to be among the class of mental actions are happenings that are not so assessable. Instances of unbidden imagery or unbidden memory, of having one's attention captured by a loud noise or flash of light, of bare perceiving, and so on are not happenings that admit of criticism or praise with regard to their practical rationality. The principles behind this way of intuitively grouping the relevant mental happenings are therefore partly illuminated by the action condition of the Neutral Account. That is, these mental happenings are non-actions because they are not responses to specifically practical reasons.

The second reason to adopt the Neutral Account for fixing the topic of discussion is that it outperforms its most commonly espoused alternatives. These alternatives claim that for some action to count as mental, it must either occur in the absence of bodily movement (Peacocke 2009; Dorsch 2012; Wu 2023a; *cf.* Levy 2019; Vierkant 2022) or must not be in principle subject to third-person observation (Soteriou 2009). Call the former the No Bodily Movement Account and the latter the Covert Account. On both accounts, the relevant contrast to mental action is not mental happenings that are not actions but, rather, bodily or observable actions, respectively. Bodily and observable actions are those consisting at minimum in intentional bodily movements. To count as a mental action, then, is to be an action that is neither bodily nor observable. Thus, to count as a mental action is to be the agent's forming a mental state for a practical reason where this means that no corresponding bodily movement or other observable event occurs.

Both the No Bodily Movement Account and Covert Account are too strong. For one, they fail to classify as mental actions cases like the following. An agent engages in a bit of mental arithmetic using an imagined abacus and, as part of this action, her hand moves as if moving the beads of the imagined abacus. Intuitively, that she moves her hand should not disqualify what she does from counting as an instance of mental arithmetic and, thus, from counting as a mental action proper. This is true even if we stipulate that the hand movement is an important part of her calculating, say, by helping affect the appropriate modifications of her abacus imagery. Similarly, in switching attention from one conversation to another, an agent might well move her eyes or head in the direction of the latter conversation. Intuitively, her switching the focus of her attention is a mental action even if it involves as a partial means the movement of her eyes or head. These movements would not occur on this occasion were it not for the mental action that she is performing, and they might well improve her perceptual contact with the object of her attention. In which case, such movement is arguably part of and made intelligible in light of her mental action. The same is true of the abacus-using mental calculator. Yet, in both cases, both the No Bodily Movement Account and the Covert Account deny that what happens with the agent's body is part of any mental action. I will have more to say in response to this denial in chapter 2.

Second, the No Bodily Movement Account and Covert Account both beg the question against extended mental action (Vierkant 2022). Consider a variation on the classic example in Clark and Chalmers's (1998: 12*ff.*) argument for the extended mind. Otto, an Alzheimer's patient, uses his notebook to remember the location of the Museum of Modern Art (MoMA) in Manhattan. Let's assume that Otto's use of the notebook satisfies the conditions on counting as extended memory that Clark and Chalmers set out (1998: 8-9). That is, i) Otto uses the notebook sufficiently frequently and with sufficient reliability, ii) his use of it is functionally

equivalent to a healthy subject's use of her memory, and iii) taking the notebook away would severely hamper Otto's ability to remember. Granting that use of the notebook thus meets the criteria for counting as an extension of Otto's mind, his intentional use of it with the aim of recalling MoMA's location or with the further aim of, say, finding his way there arguably constitutes an extended mental action of remembering MoMA's location. It would surely count as such on the Neutral Account. Yet, neither the No Bodily Movement Account nor the Covert Account can make sense of what Otto is doing in using his notebook except, at best, as a set of bodily actions that results in his coming to believe that MoMA is on 53rd Street. The possibility that Otto is engaged in a mental action by use of his notebook is precluded by both accounts. I will have more to say in response to this in chapter 2 as well.

By contrast, the Neutral Account allows that actions performed entirely without the use of one's body are sufficient to count as mental actions. But neither lack of involvement of the body or other observable happenings are a requirement on an action's being mental. Actions that come to involve some amount of bodily movement or other observable happening outside of the body as, say, sub-actions are not thereby disqualified from counting as part of mental actions. The Neutral Account is more flexible than either the No Bodily Movement Account or the Covert Account by allowing for extended mental action and aligning with intuition on cases like the mental abacus-based calculation and attention-switching cases. The account's flexibility speaks in its favor.[5]

A proponent of either the No Bodily Movement Account or the Covert Account might object with the following example. Suppose an agent wants to form the belief that the lights are on. He forms the intention to form that belief, makes his way over to a light switch, and turns the lights on. Given that the Neutral Account allows bodily actions or other observable events to count as parts of mental actions, it should classify the agent's acts of turning on the lights and subsequently forming the belief that the lights are on as constituting a single mental action. But, so the objection goes, this is an absurd result. At most what the agent does is a series of non-mental actions that culminates in the act of bringing himself to believe that the lights are on. And the fact that the whole action is one of bringing himself to believe suggests that, in fact, no mental action occurs: the belief is merely a product of the bodily actions that the agent performs.[6]

---

[5] And this is not to mention that, unlike the No Bodily Movement Account, the Neutral Account does not beg the question against embodied or ecological accounts of mind.

[6] For a description of such a case as an instance of so-called "managerial control" of belief, see Hieronymi (2006, 2009). For an argument of this kind against counting remembering as a mental action, see Mele (2009). Mele allows that bringing it

I have two responses. First, what the example really shows is that some instances of forming a belief are such that their formation is not a response to a practical reason. The Neutral Account claims that for a happening to count as a mental action it has to be one that by itself can occur for a practical reason. The mental condition secures that the doing in question is the formation of the relevant state while the action condition secures that forming that state is something the agent aims to do and, if she does it, does it for a practical reason. Together, these conditions suggest that whatever additional aim the agent might have in forming the relevant state, forming that very state serves as an at least partial means to achieving that aim. The account does not imply that a mental action can consist in things other than the formation of the relevant state that merely result in that state. In the case under consideration, the agent's practical reason plays the wrong role in the formation of his belief. That is, it is not that in virtue of which he forms the belief. Rather, his practical reason motivates him to do other things which merely bring it about that he forms the belief. But, in such a case, its formation is only wrought by the agent's being in the appropriate epistemic situation, namely, one in which he has sufficient evidence that the lights are on and which results from his turning on the lights.[7]

My second response is that, given that the Neutral Account is just meant to fix the topic of discussion and that the reasons favoring it are just meant to be *prima facie*, its possibly erring on the side of overextending the domain of mental action can be forgiven. I believe that Otto's case differs from the supposed counterexample in that, on the extended mind thesis, Otto's use of his notebook itself consists in the forming or maintaining of a mnemic state for practical reasons. By contrast, the agent's behavior in turning on the light in the supposed counterexample does not meet the criteria for extended mind and, so, does not come to constitute the formation of any extended mental states. Moreover, in line with the first response, I take it that memory

about that one remembers is a mental action. I believe his argument depends on a now dubious view of memory as a store of facts and retained experiences.

[7] By contrast, the Neutral Account treats (at least some instances of) reasoning as mental action precisely because (in the relevant instances) the actions that the agent takes towards being in the appropriate epistemic situation for terminating the relevant bit of reasoning in a judgment consist in an activity of forming the correspondingly appropriate representations where this activity is engaged in for a practical reason. One can use evidence one already has for practical reasons. Yet, one does not count as acting merely by recognizing (new) evidence. I take it that this sense of "using one's evidence" roughly corresponds with the use of "figuring out" as in "I am trying to figure out when the train arrives." And I take it that this sense of "merely recognizing (new) evidence" roughly corresponds with the use of "realization" as in "I realized that the train arrives at 3." I am not a dyed-in-the-wool ordinary language philosopher. But let's allow, following Thompson (2008), that forms of judgment can be indicative of the structure of human activity. The fact that one can utter the statement "I am trying to figure out when the train arrives," where the pragmatics of this statement imply agency and the fact that one cannot utter the statement "I am trying to realize when the train arrives" with the implication of agency together suggest that at least some instances of belief-formation are not something we do for practical reasons. I will have more to say about this in chapter 2.

differs from the relevant instances of belief-formation with respect to the ability to engage in either for practical reasons. Otto can use his notebook and a healthy neurally-intact adult human being in normal conditions can use her memory for practical reasons. While I think these responses keep the Neutral Account immune to the counterexample, the former implies that a person's mental states are not necessarily located in her skull. Indeed, as I will be arguing in chapter 2, the powers framework's account of mental action has it that many mental states are not in the skull. That said, a reader who takes this counterexample to be insurmountable by the Neutral Account or who finds that allowing for extended mental action untenable can rest easy knowing that standard causalist accounts of mental action that at least try to restrict the domain of mental action to certain events within the skull match up with the Neutral Account with respect to counting those events as mental actions.

While neither of my responses constitute knock-down arguments in favor of the Neutral Account, I take it that I have provided sufficient *prima facie* reason for preferring it. Moreover, even if a reader finds the account dubious as the ultimate account of mental action (as well they should), they should rest easy knowing that its purpose is simply to set a shared target in view for the powers framework and its causalist rivals. Both are aiming to explain in virtue of what the agent's forming a mental state for a practical reason constitutes her performance of a mental action. The aim of this dissertation is to set out the powers framework and extend it with a view to providing its explanation of mental action.

Before giving an overview of the chapters, I want to briefly set out the standard causalist account of mental action. Doing so will provide something to contrast the powers framework's account with, although I will not be assessing causalist accounts of mental action specifically. Rather, the point is merely to show that the powers framework provides a viable alternative. According to standard causalist accounts of mental action, a mental action is the formation of a mental state that is (non-deviantly)[8] caused and guided by a truth-conditional propositional directive attitude whose content represents that state as, say, to be formed by the agent. Such directive attitudes might turn out to be desires, beliefs, belief-desire pairs, intentions understood as

[8] This formulation is meant to block deviant causation. Deviant causation occurs when the causal pathway by which an agent's intention to do *A* brings about a doing of *A* that exhibits a failure of agentive control or guidance. For instance, the nephew of a wealthy man may be set to inherit his uncle's wealth and, driving one night, decides to kill his uncle. The nephew's intention so unnerves him that he loses control of the car, swerving and fatally injuring a pedestrian who just so happens to be his uncle. In this case, the nephew's intention to kill his uncle causes the nephew to kill his uncle. But, arguably, his killing his uncle is not an intentional action. This is because it is deviantly caused. For some possible solutions to causal deviance, see Wu (2016) and Shepherd (2021); however, see also Levy (2013). While I do not address deviant causation head on in the main text of this dissertation, what I have to say about how an agent relates to her action and the difference between intentional actions and lucky success in chapter 2, to my mind, suggests that causal deviance is fundamentally an issue of causalism's inability to make sense of the self-knowledge an agent possesses in acting.

primitives, plans, goal-states, and so on, depending on the specific form of causalism one adopts. They are truth-conditional in the sense that they have their content by virtue of bearing both appropriate inferential relations to other such contents and appropriate mapping relations to those objects. Their contents thus stand in for or go proxy for their objects. Such contents are "true" if the latter mapping relations pick out what the agent actually goes on to do. Call forms of causalism captured by this formulation "reductive causalism." Alternatively, one might adopt an agent-causation account of mental actions, according to which mental actions are the formation of mental states that are (non-deviantly) caused and guided by the agent herself in the capacity of a substance-as-cause.

In either case, what makes something a mental action on causalism is, at bottom, the instantiation of a causal relation between, on the one hand, the relevant truth-conditional propositional directive attitude or the agent herself and, on the other, the formation of the relevant state. Such accounts fit with the Neutral Account at least with respect to those mental actions that involve no bodily movement or are fully covert. Causalism fits with the action condition inasmuch as the relevant directive attitudes are formed in light of a practical reason or inasmuch as the agent herself instantiates the relevant causal relation in light of such a reason. Causalism fits with the mental condition inasmuch as what it counts as mental actions are a proper subset of what the Neutral Account counts as mental action. Because standard forms of reductive causalism in particular understand mental representations as truth-conditional representations somehow realized in the activity of the agent's neurocognitive architecture, such accounts have reason to restrict the domain of mental action in accordance with the No Bodily Movement Account or Covert Account. Yet, such accounts can nonetheless allow for extended mental action, assuming they accept criteria for functional parity between, on the one hand, external use of in principle intersubjectively available (or public) representations and, on the other, the computations that are supposed to be performed in light of the activity of the agent's neurocognitive architecture. In any case, the paradigmatic cases of mental action on such accounts will be purely internal "transitions in thought" constituted by such computations when those computations are (non-deviantly) caused and guided by some truth-conditional propositional directive attitude.

In treating causalism as a rival to the powers framework, I take it that the framework thinks of the relation that the agent bears to her action as not, in the first instance at least, consisting in a causal relation. Yet, on the framework, actions are nonetheless expressions of her causal powers. The framework's ability to navigate this apparent tension lies in its holding, on the one hand, that intentional actions are unified developing processes constitutive of the manifestations of habits and skills and, on the other, that the mental representation

of relevance to explaining the performance of an action is itself just an earlier stage of that process. The performance and successful completing of the action are merely later phases in a single, unified process. It is the whole process, the whole interaction between agent and the world, that is the cause of whatever the agent effects in the world. Hence why the framework holds that actions are expressions of causal powers. I will discuss these points in more detail in chapters 1 and 2. The point for now is that the framework treats action not as a type of behavioral effect with a certain type of psychological cause. Rather, it treats intentional actions as teleologically unified processes that develop in characteristic stages consisting in *a priori* "moments" of activation, ongoing performance, and completion. The ground of these moments and their unity comes from what I call the agent's "cogent practical cognizing." Such cogent practical cognizing consists in the agent's ordering into means and ends the manifestation(s) of her disposition(s) for the relevant actions in light of "wanting," in Anscombe's (1957/2000) sense, what those actions are a means to achieving, securing, or getting (§36). Such ordering can occur explicitly in either overt or covert acts of deliberation or can occur implicitly in the activation of a network of dispositions that are already calculatively ordered into means and ends thanks to some amount of training and practice. Because some degree of internal calculative order among such dispositions is required for intentional actions of the sort that human beings engage in, our intentional actions are essentially expressions of skills and, when they form an insoluble unit for the agent, habits.

What's important for now is that, in conceiving of actions as teleologically unified processes that develop in characteristic stages, the powers framework treats the agent's intention to act and her intentional action as two phases of a single metaphysical unit. Her intention is the potentiation of the intentional action, and the performance of that action is the actualization of that potential. What is potentiated in intending and actualized in performance is a causal power, a way that the agent interacts with and changes the world. So, causation has a role in the powers framework at least with respect to what an agent is acting on through her performance. But, according to the powers framework, the action itself is not an effect, at least not when considered from the point of view of the metaphysics of action and agency. That said, the framework allows that there are causal connections between the enormously complicated activity of the agent's neurocognitive architecture, her musculature, and things in the world any time the agent performs an intentional action that involves her body. But allowing this should not be confused with identifying such causal relations or the relevant activities as constitutive of action. Rather, they are relations and happenings that are given correct descriptions which do *not* "go beyond physics" (Anscombe 1957/2000: §47). They are, in a sense, matter that is given form and unity by the agent's causal powers and cogent practical cognizing by virtue of underwriting the

expression of such powers in intentional actions. I will have more to say about agential powers and the specific rational form such powers take in human beings in chapter 1. For now, I will end this introduction with a disclaimer and a brief overview of the chapters of this dissertation.

The disclaimer is that nowhere in the dissertation do I take myself to provide any knockdown arguments against causalism. I raise some concerns about causalism in chapters 2 and 4 and in Appendix F. But none of these concerns are meant to show that causalism is untenable either as an account of mental action specifically or as an account of action in general. My aim in raising these concerns is merely to make room for the powers framework as an alternative. Both by providing the positive view and by showing how it differs from causalism, I merely hope to show the viability of the powers framework as an account of mental action. So long as the reader can see how the powers framework provides a legitimate account of mental action that is not equivalent to reductive causalism or agent-causation, I have done my job.

Overview of the Chapters

In the first chapter, I give a synoptic overview of the powers framework. It holds that intentional actions in general are the expression of skills and habits and that skills and habits, in turn, are intelligent dispositions for action. One part of what makes these dispositions intelligent is that they are rationally transmitted through what Small (2017) calls the "life cycle of a skill." The other part of their intelligence is a product of such rational transmission, namely, the agent's being in a position to order the dispositions she is acquired in acquiring the relevant skill into means and ends that she thereby also recognizes as ways that any agent situated as she is can go about achieving, securing, or getting what she is after. I extend Small's account by introducing four relations among dispositions for action which I call "incorporation," "subsumption," "determination," and "integration." The life cycle of a skill can then be understood in terms of expert-guided sub-cycles of incorporation of skills by habits and subsumption of habits by skills that result in determinations of the shape of a skill and integration between skills.

I combine Small's powers-metaphysics with Springle's non-reductive action-first account of intentionality. According to this account, representations in general are a matter of what Springle calls "practical access." Mental representations consist in the potentiation of interactions between a subject and things in the world and provide *unmediated* practical access. Mental representations represent their objects as affording opportunities for the relevant type(s) of interaction whose generation by the subject would satisfy some need. Such interactions are potentiated either when the subject comes into perceptual contact with

something in the world or when the subject develops a need that would be satisfied through the relevant type of interaction with the relevant type of thing. The subject generates such an interaction when it comes into perceptual contact with the relevant type of thing, when it has the relevant need, and when such generation is not impeded by the potentiation or generation of another type of interaction inconsistent with its generation. Such interactions constitute actions.

In human beings, a significant subset of types of interaction that we can generate consist in the construction and use of objects that stand in or go proxy for other things by bearing appropriate mapping relations to those other things and by bearing appropriate inferential relations to each other. The potentiation and generation of these latter types of interaction provide *mediated* practical access to things in the world that may or may not be present and such interaction represent these things as what they are, as having certain properties, as standing in certain relations to other objects, and so on. It is here that we get truth-conditional representations. Moreover, only the actual generation of the relevant types of interaction constitute the formation of such representations. That is, a truth-conditional representation of something is formed only when an agent actually interacts appropriately with (or to construct) immediately accessible representational vehicles that stand-in or go-proxy for the thing represented. I call these "discursive representations." Because mental representations are not of themselves interactions the organism generates with representational vehicles, there are no actual discursive representations in the head on the powers framework. The only internal representations of things are those that present their objects as affording *opportunities* for interaction. In humans, a significant subset of these representations present their objects as affording opportunities for constructing and using discursive representations. When it comes to rational agency, potentiation of intentional actions comes with potentiation of acts of deliberating, expressing intention, or action description, assessment, or explanation concerning those actions. And it is in virtue of such dual potentiation that the agent represents what she is doing both as acting on an affordance and as doing something in a way that any agent situated as she is could do.

Combining Small's and Springle's account, I set out the form of practical reason and, drawing from Anscombe (1957/2000: §§33-44), identify what drives it with what I call "practical want." I identify practical reason driven by practical want with practical cognition. I claim that, when an agent calculates an order of means and ends whose instantiation would actually satisfy a need she has, her practical cognizing is cogent and constitutes an exercise of practical knowledge constitutive of a corresponding intentional action (I give a briefish overview of Anscombe's account of practical knowledge in Appendix A). I thus adopt what Schwenkler

(2019: 99) calls a "factualist understanding" of practical knowledge and, therewith, what some have called the "knowledge thesis" or "cognitive condition" on intentional action (for instance, Small 2012). The latter amounts to the claim that to act intentionally an agent must possess and be exercising a corresponding bit of practical knowledge (I defend this claim in chapters 1 and 2 as well as in Appendices B-E). Finally, I return to the life cycle of a skill and give an elaborated description of it using Springle's account and my own extensions of the powers framework. In setting this up, I touch on so-called "basic actions" and the role of propositions in the life cycle.

In chapter 2, I extend the powers framework to address mental actions. According to the framework, mental actions are expressions of mental skills and mental habits. What unifies all three as *mental* is their being intentional actions of or dispositions for, respectively, constructing and using discursive representations. Overt mental actions consist in the actual construction and use of such representations. Covert mental actions consist in the potentiation of such construction and use sufficient to potentiate further construction and use of what would have resulted from the actualization of the former potentiation. I identify the latter potentiation with the exercise of recognitional capacities. I briefly relate the powers framework's account of mental action to the life cycle of a skill and touch on the upshots this has for understanding the relationship between theoretical and practical reason. I then elaborate three dimensions of the powers framework's account of mental action. First, human perception is conceptual by virtue of being recognition-based and, thus, consists in the potentiation of a multitude of (basic) mental actions. Second, mental actions are type-individuated by the type of need such actions function to satisfy and by the corresponding type of interaction with targets sufficient to satisfy that type of need. Third, mental actions start out as extended, are often extended, and are only sometimes covert. Finally, I address two objections. First, drawing from Strawson (2003), one might object to there being mental actions consisting in the direct production or modification of content in consciousness. I discuss five possible responses to this objection, settling on the last, namely, the rejection of causalism. Second, Wu, in conversation, expressed a worry some time ago that the account of habitual mental action that I was then trying to articulate could not distinguish genuine instances of such actions from lucky accidents wherein the desired content just so happens to "pop" into the agent's head. Given that I have since undergone something of a radicalization, the worry is even more pressing and attaches to the powers framework's account of intentional action more broadly. I resist this objection by pointing out how the powers framework distinguishes intentional actions and lucky accidents in general and how it distinguishes habitual mental actions from instances where the desired content luckily and accidentally "pops" into the agent's head.

In chapter 3, I drill down on imagination, providing what is, to my knowledge, a novel account of this faculty according to which it is paradigmatically a type of skillful mental action. More specifically, to imagine is to at least potentiate the formation of a discursive representation the content of which the agent directly selects and where such (potentiated) formation and selection together constitute a means to the performance of learned behaviors such as pretense, engagement with fiction, predicting others' behavior, reasoning about possibility and necessity, hypothetical and counterfactual reasoning about contingent matters of fact, mind-wandering, imagining for its own sake, and so on. I place such behaviors together under the heading of "imaginative projects." Completion of an imaginative project like pretense, then, stands to the active (potentiated) formation of the corresponding representation as an end. The account thus bakes in the calculative, means-end order that the powers framework, following Anscombe, takes to be definitional of intentional action. Moreover, drawing from chapter 1, I hold that a central constitutive feature of this order is the agent's possessing and exercising practical knowledge both of her act of imagining and of its having the relevant means-end order. Her practical knowledge constitutes her corresponding act of imagining.

I argue that the powers framework's account of imagining has the potential to unify the seemingly disparate suite of imaginative projects. In particular, characterizing imagining as paradigmatically a type of skillful action can explicate its role as a means in contributing to pretense, engagement with fiction, and so on. According to the account of imagining I offer, like other types of skillful action, distinct acts of imagining can be put to use towards very different and even conflicting ends. Moreover, in their service as means, such acts can take on very different, even conflicting properties. For instance, painting as part of creating a work of art might well require the agent to do something very different or that even conflicts with what would be required of her in painting as part of renovating a home. Nonetheless, painting as a type of skillful action is identifiable despite this variation and conflict across instances. If imagining is paradigmatically a skillful mental action, then, such variation and conflict across its instances is to be expected and is nonetheless consistent with type-individuating it along the lines suggested in chapter 2.

I then argue that the account makes better sense of the standard view of the intentionality of imagination. According to the standard view, an agent imagines what she intends to imagine. On the powers framework's account, given the possibility of various kinds of *prima facie* failures of practical cognition (discussed in Appendix C), an agent can fail to imagine in accordance with her aim and, in such cases, she misimagines. Finally, I address seemingly passive instances of imagining like unbidden imagery, automatic and involuntary imagining, and mind-wandering. I argue that unbidden imagery is likely not best thought of as

instances of imagining, that automatic imagining is consistent with imagining being skillful, that involuntary imagining can be understood in terms of misfires or slips in the manifestation of the capacity to imagine, and that mind-wandering can be understood as automatic or as an action slip. I conclude with six avenues for further research.

Finally, in chapter 4, I show that, given the powers framework's understanding of mental action and imagination, episodic remembering turns out to be a subtype of imagining by virtue of being a kind of imaginative project. More specifically, episodic remembering is a subtype of episodic thought more generally, the kind of thought involved in episodic hypothetical or counterfactual reasoning about contingent matters of fact. Episodic remembering, then, involves an agent's at least potentiating the construction of some scene and selecting its content with the aim of (veridically) representing some specific event from her personal past such that she re-experiences that event. In work with Springle (under review) I provide a more complete account of episodic memory as a kind of performative "historytelling." In chapter 4, however, I only seek to show that the powers framework gives us resources for arguing that episodic remembering is a kind of imagining on the assumption that such remembering is a kind of mental action.

I draw out two commitments of accounts of remembering as a mental action. First, episodic remembering is relevantly like imagining in virtue of being constructive. Second, the construction involved in such remembering is highly constrained by how things were in the agent's past. These two commitments seem to be in tension with each other. And the tension is made worse by identifying both episodic memory's constructive and constrained aspects with the activity of mechanisms in the agent's neurocognitive architecture. I argue that the tension can be resolved by thinking of both of these aspects in more agential terms. That episodic remembering consists in the (potentiation of the) construction of scenes and that the agent has at least some degree of freedom with respect to the content of such scenes suggests that she is involved in the selection of content in the same way that she is when she is imagining. Nonetheless, the constraints involved in episodic remembering mean that she cannot construct "whatever she wants" and still count as remembering or trying to remember. I claim that such constraints are self-imposed and satisfied by the agent's action. Indeed, all action is constrained by the agent to some degree. I call the agent's imposing constraints and satisfying them in action her "self-determination," since it provides a lens through which we can appreciate the agent's actively structuring her own agency (I consider such active structuring in discussing practical cognition in chapter 1 but without appeal to constraints). And all imagining likewise involves the self-imposition and satisfaction of constraints. Epistemically useful imagining, then, involves the self-imposition and satisfaction of specifically epistemic

constraints. And episodic memory is continuous with such epistemically useful imagining. I conclude by considering whether extant accounts of remembering as a mental action can adopt my self-determination-based resolution of the seeming tension between episodic memory's constructive and constrained aspects.

1.0 The Powers Framework

1.1 Introduction

In this chapter, I introduce and extend the powers framework for action in general. I do this for two reasons. First, the account of mental action, mental habit, and mental skill to be given in the following chapter is itself an extension of the framework set out here. Since this framework contrasts with causalism of the sort that is prevalent in the contemporary literature on mental action, it helps to have a grasp of the underlying metaphysics independently of the account to be defended in this dissertation. Second, in laying out the framework, I extend it in ways that anticipate later chapters. Specifically, I introduce relations between the types of dispositions for action that others, specifically Small (2021), have set out in constructing the framework. The explication of these relations helps to show how rational agency consists in the exercise of complex or "multitrack" dispositions. In furtherance of this goal, I tie the connections between and within these dispositions to the representational contents that are their activation and manifestations. In doing so, I aim to set out in rough outline the domain of mental action as exercises of rational capacities for the construction and use of truth-conditional or what I will later call "discursive" representations. By contrast, non-mental actions performed by human beings consist in the exercise of rational capacities for other sorts of endeavors whose exercise is, at least in part, expressions of judgments that are in turn primed for discursive representation. Or so I will try to show in this chapter.

Section 1.2 sets out the metaphysics of the powers framework. According to the framework, skills and habits are dispositions for intentional actions that are internally calculatively structured and subject to norms specific to the relevant domain or practice. Section 1.3 sets out the intentionality of action, what I will call "the *generans*" of practical reason, and the form of practical reason. At bottom, the intentionality of intentional action consists in practical knowledge. The exercise of practical knowledge in turn consists in the cogent drawing on of intentional action concepts. And such drawing on in turn consists in the instantiation of the calculative means-end order in action that those concepts represent. The intentionality of skills consists in the possession of such practical knowledge in the form of knowledge-how to perform the relevant actions where that possession is realized both in the internal structure of skills and in the inter- and intra-dispositional connections between and within them. Finally, the intentionality of habit consists in the possession of practical knowledge in the form of knowledge-to perform actions where that possession is realized in the internal structure of habits and in the inter- and intra-dispositional connections between and within them. Section 1.4

ties the threads of Sections 1.2 and 1.3 together. The relations within and between rational dispositions for action set out in Section 1.2 depend on the form and "*generans*" of practical reason as set out in Section 1.3. The agent considered as agent, then, is the unity of her dispositions for rational action, which dispositions themselves are organized complexes of domain-specific intentional action concepts organized into calculative means-ends relations and domain-specific norm-driven practical-inferential relations. The impact of the development of capacities for constructing and using discursive content on the rest of the system of dispositions is transformative. An agent capable of rationality is one whose unity as an agent is made possible by being a possible object of her own active practical reasoning, in the first instance, through her skills and habits being possible objects of her own practical reasoning. Realizing this paves the way for an account of mental agency based in the powers framework.

1.2 The Nature of Habit and Skill

To my knowledge, the metaphysics of the powers framework is most thoroughly developed by Will Small over the course of several papers and chapters. So, I will focus primarily on his work throughout this section and the next. Such work draws inspiration from two threads in the philosophical literature on intelligent action. The thread that will occupy me in this section originates in the work of Ryle (1949) and is concerned with the metaphysical structure of dispositions specific to intelligent agents. Intelligent capacities are Ryle's primary target, skill being an exemplary species of the genus. A skill, say, for playing tennis is within the same broad category as the fragility of a vase, the solubility of sugar, and so on. But unlike these latter dispositions, which manifest in a single type of event, namely, shattering and dissolving, skills manifest in a variety of ways. The tennis pro manifests her skill for playing tennis in making a serve, returning a volley, noticing the trajectory of the ball mid-play, practicing her strokes, teaching and coaching another to play, and so on. Accounting for the internal structure of a disposition like the capacity to play tennis at the professional level is what occupies Ryle and the thread in the literature that he initiates. Philosophers have recently started investigating the internal structure of habits as dispositions for action distinct from skill (for instance, see Bermúdez and Felletti 2021). While Ryle uses "mere habit" as a foil to skill, namely, as a kind of disposition that manifests in just one way, habits too seem to differ structurally from dispositions like fragility, solubility, elasticity, and so on. In particular, while habits *can* become such that they manifest in just one way, they need not always "atrophy" due to lack of use or over-practice. The present segment of this thread of the philosophical literature is concerned, then, with explicating the internal structure of both skills and habits and the relationship, if any, between them.

The second thread informing Small's framework originates in the work of Anscombe (1957/2000) and is concerned with the nature and intentionality of intelligent agency writ large. It will occupy me in the following sections.

1.2.1 The Forms of Disposition

Habits and skills are dispositions for action. Their being dispositions for action distinguishes them from dispositions like fragility and solubility as well as from functional biological dispositions like that for the release of histamines in response to infection by a foreign body. While Ryle (1949) attempts to distinguish skills further form *mere* habits by appeal to the former's being what he calls a "multitrack disposition" and one whose manifestation can be intelligent or unintelligent, we'll need a more thoroughgoing taxonomy to give the powers framework a fighting chance as an account of action. Small (2021: 240) provides a sketch of such a taxonomy in terms of the distinct forms that dispositions can take (Figure 1.1):

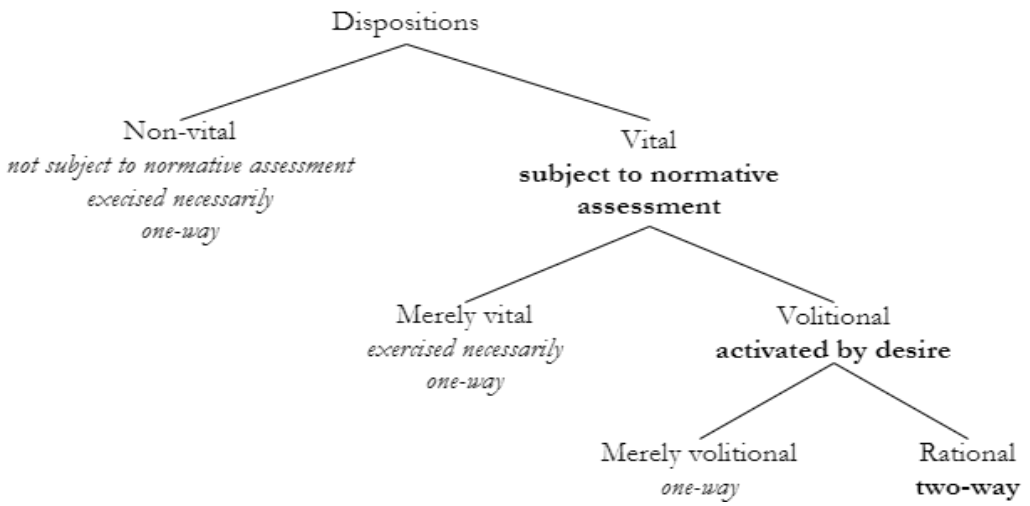

Figure 1.1 Forms of Disposition
A tree diagram depicting types of dispositions distinguished by their form. Adapted from Small (2021).

Small does not elaborate what the form of a disposition is except by way of distinguishing sets of dispositions by their form and with examples. I will not give an analysis of form for dispositions in general. However, I will give a more thoroughgoing explication of the forms of volitional and rational dispositions and of how these dispositions interact in the acquisition and mastery of skills as well as in the development of more or less intelligent habits.

The following are the forms of disposition as Small (2021) understands them. Non-vital dispositions are those most familiar to the discussion of dispositions in analytic metaphysics: fragility, solubility, elasticity, and so on. These dispositions manifest in a lawlike way when, other things being equal, the triggering conditions for their manifestation are met. Because non-vital dispositions are the causal powers of non-living things, their manifestation is not subject to assessment against normative standards internal to them. The manifestation of a physical or chemical disposition is not by itself better or worse however complete or incomplete it is or however seamless it is. Rather, such manifestations are only assessable against standards imposed by something external to the relevant dispositions. For instance, the dissolving of some substance in stomach acids is better or worse for the animal that is digesting. It is not better or worse relative to the activity of the acids independently of the animal. Rather, the vital disposition to digest imposes a normative standard from outside onto the manifestation of the non-vital disposition of stomach acids to dissolve substances whenever such a manifestation comes to partially constitute an instance of digestion.

Vital dispositions are those most familiar to discussions of biological normativity and function in the philosophy of biology. They include the disposition of histamine release upon detection of a foreign body, the disposition to digest upon ingestion, the disposition of positive heliotropic plants to bend towards the sun, and so on. Like non-vital dispositions, vital dispositions manifest in a lawlike way when, other things being equal, the triggering conditions for their manifestation are met. However, unlike non-vital dispositions, the manifestation of a vital disposition is subject to assessment against normative standards internal to that disposition. A positive heliotropic plant's bending towards the sun can be better or worse. The internal normative dimension of a vital disposition's manifestation is tied to the interdependence between that manifestation and the manifestation of other vital dispositions that together constitute a biological system's maintaining itself or reproducing in the way that biological systems of the relevant type (ought to or normally) maintain themselves and reproduce. It is because bending towards the sun contributes to the positive heliotropic plant's self-maintenance or reproduction in the way that such bending in positive heliotropic plants in general (ought to or normally) so contributes that the manifestation of the disposition for such bending can be judged better or worse. That is, such a judgment is made in reference to how such bending contributes to self-maintenance or reproduction of an individual relative to its kind.[9] Finally, while vital dispositions are causal

[9] This conception of biological normativity is at odds with several accounts of biological function. Some philosophers deny that any biological phenomena are intrinsically normative (Cummins 1975). Others allow for intrinsic normativity of the biological but hold that this normativity is grounded in statistical typicality rather than anything metaphysically lawlike (Boorse 1977). Yet other philosophers allow intrinsic normativity and might allow that there are laws of nature in biology

powers of living things, whether they manifest or not is not something that their bearers, considered as unified individuals, have a role in. One way to put this point is that a complete explanation of positive heliotropism in plants need not refer to the whole plant as doing anything which triggers or guides its bending towards the sun. One need only appeal to the relevant cellular and molecular mechanisms within the plant and the relations those mechanisms bear to the sun's ultraviolet radiation. By contrast, a disposition of a living thing whose manifestation does in some way involve their bearers, again, considered as unified individuals, is either a volitional or rational disposition.

Volitional dispositions are those most familiar to discussions of habits and animal action. They include the disposition to look for a light switch upon entering a darkened room, the disposition to take one's regular route home upon leaving work, the disposition to brush one's teeth upon waking up or heading to bed, the disposition to incessantly check one's smartphone for seemingly no reason, and so on. Like vital dispositions, the manifestation of a volitional disposition is subject to assessment against normative standards internal to that disposition. The search for a light switch can succeed or fail or be better or worse, as can an attempt to check one's smartphone, say, if the phone's battery has died or if one fumbles when trying to get it out of their pocket. A routine drive home can fail in the event of a crash or can be better or worse with respect to, say, efficiency or safety. Likewise, a teeth-brushing can fail if there is no toothpaste around or can be better or worse with respect

---

but nonetheless deny that the actual dispositions of any living thing are intrinsically normative (Neander 2017; *cf.* Millikan 1984). To my knowledge, only one family of accounts of function allow that the actual dispositions of individual organisms are intrinsically normative and lawlike in their manifestation—organizational accounts (Mossio et al. 2009; *cf.* Millikan 1984). However, for arguments against these accounts, see Garson (2017) and my (2023). One might object that this is a non-issue, since vital dispositions are best identified with what Thompson (2008) calls "natural historical judgments" and "life forms" rather than with a notion of function. Natural historical judgments are true generics that capture the manifestation of the corresponding vital disposition, for instance, "positive heliotropic plants bend towards sunlight." The life form of an organism is the collection of facts picked out by natural historical judgments specific to its kind. Yet, the move to natural historical judgments and life forms threatens to undercut the appeal to biological dispositions as the ground floor of intrinsic normativity. As I see it, the idea is to build up to rational capacities in a way that legitimizes their normativity. Presumably, this means, among other things, that the normativity of rational capacities is either itself or ultimately stems from a form of normativity that is natural in virtue of being biological. If there is no scientifically respectable naturalized normativity on offer, the normativity of rational capacities threatens to float free from the natural world. And, as a rule of thumb, we are right to defer to the biological sciences and philosophers of biology in any attempt to naturalize the normative at the level of biological processes. The same heuristic would apply with respect to the mind sciences and empirically informed philosophers of mind if we were attempting to naturalize the normative at the level of psychological processes. None of this is meant to suggest that either the biological sciences or the mind sciences is the final arbiter in what counts as a scientifically respectable naturalized normativity. Nor is meant to suggest that philosophers of biology or empirically informed philosophers of mind have the final word on the matter. It is just to say that much more work needs to be done to connect the project of naturalizing normativity to the project of legitimizing the powers framework.

to efficiency or conduciveness to gum health. As these examples suggest, the internal normative dimension of a volitional disposition's manifestation is in the first instance tied to standards of success or appropriateness in the performance of an action rather than to the contribution of that manifestation to the organism's self-maintenance or reproduction (I will have more to say on appropriateness in Section 1.3.1).[10] That the manifestation of volitional dispositions are assessable relative to standards of success or appropriateness is due to the fact that, unlike non-vital and vital dispositions, such manifestation requires, in addition to the appropriate opportunity or triggering stimulus, a desire or need on the part of the agent which that manifestation functions to satisfy.[11] The agent wants to better see things in the room, wants to get home, wants to keep her routine going or keep her teeth healthy and clean, or wants to apprise herself of the latest TikTok gossip. This additional agential aspect of manifesting a volitional disposition also means that such manifestation is not lawlike.[12]

All that said, agents do not play a role as agents when it comes to the quality of the manifestation of their volitional dispositions. That is, the agent does not determine whether the manifestation of one of her volitional dispositions is better or worse in the sense that she cannot decide to manifest its contrary. She cannot decide to manifest a volitional disposition of hers in a way contrary to its paradigmatic manifestation. For instance, assuming, say, driving her normal route home and turning on the light are habits of hers, her intentionally driving that route haphazardly or walking around a pitch-black room without looking for a light switch are not expressions of those habits. This is just what it means for volitional dispositions to be "one-way," as per Figure 1.1. If she could determine whether the manifestation of the relevant disposition was better or worse such that she could manifest that disposition in doing what is contrary to its paradigmatic manifestation then that disposition would not be volitional—it would be rational.

This brings us, finally, to rational dispositions. Rational dispositions are those most familiar to discussions of intelligent action and skill. Most importantly, *rational dispositions include our capacities for intentional action.* That said, these dispositions are arguably most fully on display in the exercise of capacities we tend to identify as "expert skills," for instance, the capacity to play tennis at the professional level (Montero

[10] I say "in the first instance" because the manifestation of volitional and rational dispositions can also be assessed along ethical and aesthetic dimensions. And there is no reason against thinking that their manifestation can be assessed against whether or not they contribute to survival or health, rearing children, and so on.

[11] I am borrowing the language of needs and a manifestation's having a function to satisfy such needs from Springle's work on action and mental representation. I will return to her account in the following section.

[12] It might be that non-vital and vital dispositions are also not truly lawlike. But it would be especially odd to claim that the law that holds in the case of manifesting the disposition to turn on the light implicitly includes something like "unless the agent decides not to" in its *ceteris paribus* clause.

2016, Foultier 2022). Like vital and volitional dispositions, the manifestation of a rational disposition is subject to assessment against normative standards internal to that disposition. For instance, a serve in a game of tennis can succeed or fail or be better or worse across dimensions of assessment tied to the domain or practice of playing tennis. Like volitional dispositions, the normativity internal to rational dispositions is in the first instance tied to the success or appropriateness of performance rather than to any contribution such a performance might make to the self-maintenance or reproduction of the organism performing the action. And, like volitional dispositions, this difference from vital dispositions is due to the fact that manifesting a rational disposition requires, in addition to the appropriate opportunity or triggering stimulus, that the agent has a desire or need that manifestation of the rational disposition functions to satisfy.

However, unlike volitional dispositions, the agent manifesting a rational disposition does play a role in determining whether that manifestation is better or worse in the sense that acquisition of that disposition implies an ability to intentionally manifest its contrary. For instance, a pro tennis player can intentionally fail to make a serve in any number of ways as part of demonstrative instruction for novices (Small 2021). That agents can manifest their rational dispositions in ways contrary to those dispositions' paradigmatic manifestations is just what it means for them to be "two-way" rather than one-way.

This last point extends beyond instructive intentional failure. Consider a seasoned doctor in internal medicine. His extensive knowledge of treatment allows him to discern, say, the appropriate dosage of morphine for any given patient. Yet, this knowledge likewise allows him to intentionally overdose a patient. Following Aristotle, in playing the role of poisoner, the doctor exercises his medical expertise such that "by denial and removal [...] [the manifestation of a rational disposition] exhibits the contrary; for the contrary is the primary privation, and this is the removal of the positive term" (*Metaphysics* Θ.2 1046b14-15; see also Small 2021: 239-243). Small claims that the ability of an agent in possession of a rational disposition to exercise it both in expressing its positive term as well as in removing that term is due to the fact that "a rational disposition consists in *productive knowledge* of the normative layout of the domain of its manifestation—knowledge, that is, of the good and the bad, as they are determined by the standards internal to the disposition" (2021: 241; original emphasis). In learning a skill, one's goal is to come to know the good of the domain of its manifestation. In the case under consideration, this consists in knowledge-how to medicate. But such knowledge-how will bring with it some knowledge of the bad for free, namely, knowledge-how *not* to do what is considered good by the standards internal to the disposition. This is just what the case of the poisoning doctor shows. Finally, the same insights apply to a rational disposition that stands contrary to another. Consider as a contrary to the art of

medicine the art of poisoning, granting the artificial restriction in the scope of medicine that this implies. In acquiring expertise at poisoning, an agent comes to know some things about how to medicate, where that knowledge consists in, for some specific poisons, knowing how *not* to successfully poison.[13]

So much for the forms of disposition. Small's account goes beyond this set of distinctions. He complicates his taxonomy by arguing that both volitional and rational dispositions in human beings exist on a spectrum of intelligence. The addition of a spectrum of intelligence for volitional and rational dispositions makes for a fourfold distinction. On the one hand, volitional dispositions range from what I will call unintelligent "rote routines" to intelligent tendencies. On the other hand, rational dispositions range from mere competences to what I will call "expertise."[14] In addition to the introduction of a spectrum of intelligence for volitional and rational dispositions, elsewhere, Small presents a metaphysics of education that implies a kind of particularism concerning (at least) rational dispositions. However, before turning to consider these complications, I add one of my own, namely, laying out four relations within and between volitional and rational dispositions. Hereafter, I will refer to rational dispositions as skills and volitional dispositions as habits, except where clarity dictates.

1.2.2 Four Relations Within and Between Volitional and Rational Dispositions

Habits and skills in human beings are not just dispositions for action that exist independently of each other. Nor are they simple dispositions consisting of a single causal power. Both are complexes of rationally interrelated causal powers and each can come to contain the other. Moreover, their internal structure and the relations between them are born of the intentionality of their constituent powers. In Sections 1.3 and 1.4, I will unpack what all this consists in. For the time being, I will summarize: the internal structure of both habits and skills consists in intentional action concepts specific to the relevant domain that form networks of calculatively

[13] In neither case does this knowledge-how to do badly in one domain require understanding the mechanisms that, arguably, one would have to know to count as knowing how to do well in the contrary domain. The doctor ought ideally to know how morphine is metabolized by the body. Knowledge of the mechanism can itself give him insight into how overdosing is efficacious at ending a patient's life. Nonetheless, he might not know how exactly morphine kills and still exercise his expertise in poisoning a patient with it. Similarly, an expert poisoner ought ideally to know how, say, sweet wormwood is metabolized such that it kills. Nonetheless, she might not know how exactly ingesting small amounts of sweet wormwood is beneficial and still exercise her expertise in medicating someone she does not wish to kill.

[14] I change the terminology Small (2021) uses for ease of reading. He uses "rote habit" and "mature skill" where I use "route routine" and "expertise" respectively. Because I am identifying rational dispositions with skills and volitional dispositions with habits, I want to avoid confusion in the use of terms. Mere competences are still skills. And intelligent tendencies are still habits.

structured means and ends. Habits and skills are intelligent in virtue of these networks of calculatively ordered intentional action concepts and in virtue of the agent's ability to deploy (parts of) these networks so as to satisfy her desires or needs. Moreover, these networks are related to each other in a number of ways, at least one of which reflects a further dimension of their intentionality. Namely, networks of calculatively ordered intentional action concepts are related to each other both within and across habits and skills by relations of weak practical consistency. Two intentional actions are weakly practically consistent if they can both be successfully performed intentionally at roughly the same time. An agent's weakly consistently intending to do both amounts to the activation of corresponding rational dispositions the full manifestations of which are both possible for her at the time of activation. I will have more to say about intentional action concepts, the calculative means-end order, and practical consistency in the following sections and in Appendix E. In this subsection, I introduce four relations within and among habits and skills. I call them *incorporation*, *subsumption*, *determination*, and *integration*, respectively. I elaborate these relations and their mechanisms in the following sections and subsections. For now, I want to lay them out and draw attention to some of the upshots for habits and skills being so related.

Let's start with incorporation:

Incorporation: a habit incorporates a skill just in case the manifestation of that habit is at the same time a manifestation of the relevant skill.

For instance, manifestation of the habit of driving one's normal route home from work is always at the same time the performance of an act of driving. The capacity to drive is a rational disposition, a skill. Among its manifestation conditions are desires or needs that skill for driving functions to satisfy, in this case, a desire or need to get home from work and a desire or need to do so by car. One can succeed or fail to manifest the skill, again, say, in the event of a crash or can manifest the skill better or worse, say, if one unwittingly misses a turn on the way. Most importantly, in possessing a skill for driving, one is in a position to manifest that capacity in a way contrary to its paradigmatic manifestations. That is, one can intentionally make mistakes in instructing another by demonstrating how *not* to drive or one can intentionally aim, say, to hit pedestrians or drive off of a cliff. On the powers framework, so-called "basic actions" will turn out to be manifestations of habits that incorporate some skill (see Section 1.4). In which case, some degree of incorporation is a requirement on the possibility of intelligent agency. What matters for my current purpose is to recognize that habitual actions (and basic actions) always involve the manifestation of some skill(s) as well. Every habit incorporates some corresponding skill.

Let's move on to subsumption:

> **Subsumption:** a skill subsumes a habit just in case the manifestation of that skill is at the same time a manifestation of the relevant habit.

For instance, some instances of manifesting the skill for playing tennis, especially in cases where the player is a pro, are at the same time manifestations of the habit, say, of slicing thus-and-so. Slicing thus-and-so is a habit for the pro just in case she has developed a habit to so slice and intentionally not doing so or intentionally doing so incorrectly would require her to resist performing the slice (more on this in the following Sections). Encountering and having to resist the initial inertia in activation of the habit to slice thus-and-so when attempting to avoid the act or to perform its contrary suggest that it is *not* a two-way disposition. Yet, it is a disposition that has among its manifestation conditions a desire or need that the disposition normally functions to satisfy, say, to return a certain kind of serve or volley. Moreover, the tennis pro can succeed or fail to make the relevant slice or can do so well or poorly, depending on the circumstances of her action and relative to the domain of her skill. Like incorporation, on the powers framework, some degree of subsumption is a requirement for the possibility of intelligent agency. That it is a requirement follows from claiming, on the one hand, that intentional actions are the manifestations of skills and, on the other, that basic actions require manifestations of incorporating habits (see Section 1.4.2). What matters for my current purpose is to recognize that skillful actions always involve the manifestation of habit(s) as well. Sometimes this is because the skill in question is incorporated into a corresponding habit. Other times, it is because the skill in question subsumes a habit that incorporates a distinct sub-skill.

Let's turn to determination:

> **Determination:** one skill determines another skill just in case the agent possesses the first skill and the manifestation of that skill is a determinate of the manifestation of the second skill.

For instance, assuming the agent has the capacity to run a 5-minute-mile, that capacity is a determination of the capacity to run a mile, which, in turn, is a determination of the capacity to run. Running a 5-minute-mile is a way of running a mile and running a mile is a way of running. The first act-type is a determinate of the second and the second is a determinate of the third. The corresponding capacity to run is made more determinate by the agent's also possessing the capacity to run a 5-minute-mile.

Indeed, following Locke (1974: 187) and Small (2017: 255), I hold that we ought to type-individuate practical capacities—of which skills are a subset—by appeal to the "features of the [individual] agent's constitution and background which bring it about that he standardly succeeds in" manifesting the relevant

capacity, where these features do not also underwrite some more general capacity he has. If so, then we should not individuate an agent's capacity to run a 5-minute-mile as distinct from her capacity to run since the features of her constitution and background that underwrite the former capacity are those that underwrite the latter and the latter is more general. Arguably, the capacity to run is the most general practical capacity in the relevant class and serves as the highest determinable of that class. Other practical capacities within that class, then, are determinates of the capacity to run.[15] On the powers framework, it will turn out that possession of a skill always involves that skill's having some degree of determinacy beyond its highest determinable (see Section 1.4). What matters for my current purpose is to recognize that the exercise of determinate skills is also at the same time exercises of their determinables. More specific skills based in the same features of our constitution and background as more generic ones determine those more generic skills.

Finally, let us consider integration:

> **Integration:** a skill integrates a habit or a different skill just in case, in coming to acquire the first skill, the agent gives the habit or second skill ampliative application.

Two examples. First, consider a fledgling pianist who has previously developed a habit to maintain good posture when sitting at her computer. In acquiring a skill for playing, she gives this habit ampliative application: she manifests it when sitting down to type and, now, when sitting down to play. Her newly formed skill for playing integrates her habit to sit with her back straight, her feet flat on the floor, her arms close to her body and elbows slightly higher than the keys, and so on. Second, consider a fledgling pianist who has not developed a habit to maintain good posture when sitting at his computer but who knows how to do so and can do so when instructed or when he intends to so sit. In acquiring the skill for playing, he gives his admittedly mundane skill for sitting with the right posture ampliative application. His newly acquired skill integrates his skill for sitting with correct posture at his computer, even if he forgets to do so when he sits down to write his dissertation.

Importantly, the integration of some habits or skills into another skill can at the same time be an instance of determination. For example, a proficient pianist will integrate some of her habits and skills for

[15] The capacity to locomote is arguably a genus of which the capacities to crawl, walk, run, swim, climb, drive, fly and so on are species. The latter capacities involve the possession of features of the agent's constitution and background that introduce distinct modes of locomoting that are logically independent of movement. Just as rectilinearity and consisting of three straight lines are logically distinct from each other and brought together in the species "triangle," a capacity to move from one location to another is logically distinct from a capacity to move one's feet and legs rapidly enough in a cyclical motion to have a gait characterized by a short flight phase (i.e. neither foot touching the ground) and are brought together in the species "running." One can run in place (no change in location) and one can locomote by other means (no running). The introduction of a logically independent differentiating property is definitional of the genus-species relation and not necessarily present in the determinable-determinate relation.

playing piano into her newly formed skill for playing harpsichord at acquisition. At the same time, her acquiring this new skill will be a determination of her skill for playing keyboard instruments as well as her skill for playing chordophones. And the determinates of these more general skills will include those skills for playing piano that she brings to bear on playing harpsichord. Likewise, these more general skills contain the habits subsumed first by her skill for playing piano and then by her skill for playing harpsichord. On the powers framework, it will turn out that integration is a requirement for the possibility of the acquisition of skills through rational transmission. At the outset, such integration transforms arational dispositions for behavior or unintelligent action into an agent's initial battery of intentional action concepts. Acquisition and mastery of expert skills, then, is continuous with an agent's first learning to act intentionally insofar as the former likewise involves integration. What matters for my current purpose is to recognize that skill acquisition involves ampliative application of habits and, usually, other skills. Such ampliative application is critical for skill acquisition and mastery. It is to discussing such acquisition and mastery that I now turn.

1.2.3 The Life Cycle of a Skill

According to Small, skills are rational dispositions, where such dispositions are complexes of interrelated causal powers of the individuals who possess them. Small's metaphysics of education regarding skills is processual and developmental. The acquisition and persistence of an individual's skill over time mirrors the development and persistence of individual organisms. Individual organisms are produced by other individuals of the same kind through reproduction. Individual organisms maintain themselves through highly constrained dynamic interactions with the environment including the absorption of nutrients, the maintenance of a boundary between themselves and the environment, the avoidance of noxious stimuli, and so on. According to Small, self-maintenance and reproduction are distinct "moments" constituting the *a priori* unity of the life cycle of an organism (2014: 103). Skills are likewise produced in individuals (usually) by other individuals who possess the same skill.[16] Skills require maintenance as well, lest they atrophy into an increasingly discordant heap of mere competences and rote routines.

However, unlike organisms, skills are not just self-maintaining and reproductive but are also fundamentally *productive*. That is, the primary manifestation of a skill is the creation of a product through performance by the skilled individual. Such production just is the manifestation and expression of the productive knowledge in which the relevant skill consists. That is, the product of skill is the realization of the

[16] I use "usually" here to account for cases where skills are self-taught.

good or the bad of the relevant domain. For example, the primary manifestation of an expert potter's skill is the production of pottery, not practice throwing on the wheel or guiding or teaching others. Her productive knowledge just is the ability to produce at least minimally aesthetically adequate pieces or to intentionally fail to do so, say, for instructive purposes. These products, in turn, are instances of the good or the bad of the domain, respectively. The fundamental productivity of skill means that in addition to moments of self-maintenance and reproduction, the life cycle of a skill includes a third moment: doing the thing one is skilled at (Small 2014: 102-105).

Small's modeling of the development of skills on analogy with the development of organisms is meant to be a bulwark against an overemphasis on performance in the contemporary literature on skill. Indeed, philosophical discussion of skillful action is for the most part dominated by a focus on its primary manifestation, namely, expert performance. The focus on performance seems to stem from the fact that, when expressive of mastery, such performances are, paradoxically, intelligent and unreflective. The expert expresses mastery by acting appropriately in the moment or in a state of flow. Her behavior is more flexible and fine-grained than the novice's and, yet, unlike the novice, she does not need to stop and reflect about what she is doing as she is doing it (for characterizations of expert performances along these lines, see, for instance, Fridland 2014, 2017a,b, 2019, 2021). Despite holding that expert performance is skill's primary manifestation, Small contends that focusing too much on it threatens to obfuscate the nature of skill as something that exists outside of the context of performance. Indeed, a substantial part of the process of acquiring a skill occurs outside of that context (usually) through the guidance and instruction of others. What's more, an expert's continued possession of her skill depends on her putting in the hours practicing, reflecting, fine-tuning, and so on. These activities are crucial for maintaining her skill but occur outside of the context of performance. And while more mundane skills can be short-lived or acquired in a single instance of learning or practice or through the combination of already possessed skills, these more mundane skills, too, have an existence outside of performance. No human agent acquires skills in the complete absence of both practice and teaching from another.

Focus on the analogy with organisms is supposed to help us see "the three formally distinct moments that necessarily and *a priori* constitute *the life cycle of a skill: doing, practicing, and teaching*" (Small 2014: 103; original emphasis, see Figure 1.2). That is, any skill is what it is, necessarily, by being a capacity for action that is transmitted and acquired through the occurrence of these three formally distinct moments.

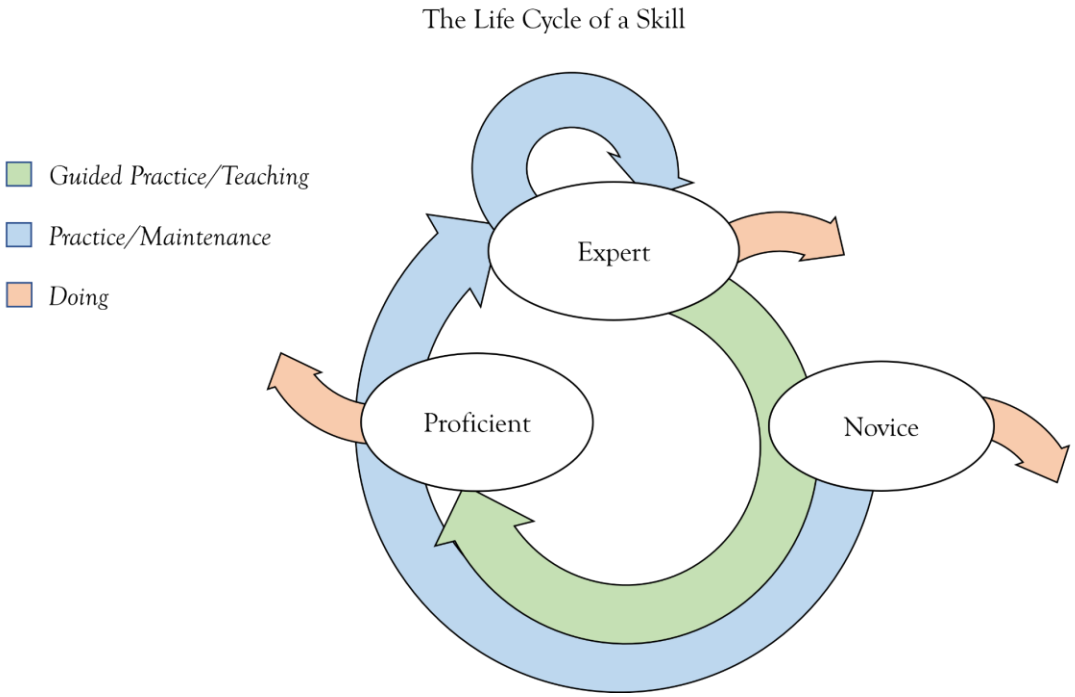

Figure 1.2 The Life Cycle of a Skill

A diagram depicting the life cycle of a particular skill as possessed by an individual. Each arrow represents manifestations of the skill in the distinct moments that together constitute the a priori unity of possessing that skill (Small 2014: 104). Each oval represents a stage in the life cycle. This diagram only represents the development of a particular skill within an individual, given that the formation of the relevant disposition in that individual (usually) depends on being guided in practice by at least one other individual.

Let's dig into the life cycle. The cycle typically starts with an expert who induces novicehood in another through guiding that individual in practice. Such practice often includes explicit instruction concerning the principles, methods, techniques, heuristics, and so on internal to the domain of the relevant skill as established by what Thompson (2008) calls a "practice" and what I will here call a "social praxis."[17,18] Social praxes unify the actions that are their exemplifications across individual performances and practitioners, act as standards for the normative assessment of those actions, and act as *explanans* for those actions as internal to the relevant praxis (Thompson 2008: 199). Social praxes, then, are sociohistorical products that allow for the rational transmission of skills. The newly initiated novice engages in manifesting the relevant skill itself, however poorly. He also

[17] Again, this ignores cases of self-teaching. The formation of new domains of skill might require unaided self-teaching. And some can acquire a skill by being purely self-taught, that is, without the aid of any kind of instruction or guidance. But both types of case are exceptional. More often, an agent who is self-taught leans on the teachings of others where such teachings have been recorded or codified.

[18] I am using the term "social praxis" to highlight both the communal dimensions of such practices as well as the rich intellectual history behind such practices. I am also using this term to allude to the Greek notion of "*πρᾶξις*," since this notion is arguably what Thompson is drawing from in his own work.

practices manifesting it and its constituent activities, at times under the expert's guidance and at times independently of that guidance. He engages in performance as well as both guided and unguided practice until he reaches proficiency. Proficiency is marked by an increase in the balance of unguided practice, improvement in the quality of performance, and increases in both the flexibility and fineness of grain of aspects of performance. Eventually, with enough practice and performance, the proficient agent achieves mastery. Unguided practice then continues to serve the function of maintaining the agent's expertise. Moreover, upon achieving mastery and assuming he has a modicum of pedagogical skill as well, the expert becomes well positioned to pass on the skill, starting the life cycle over (Small 2021: 91, fn.12).

The life cycle model is important for understanding, on the one hand, that a skill is specific and particular to its possessor and, on the other hand, that each moment in the *a priori* unity of possessing a skill is indebted to the existence of a corresponding social praxis. Starting with particularity, if Small is right that the appropriate model for skills is on analogy with the life cycle of individual organisms then, arguably, the proper level for theorizing about skill is neither the level of individual performances nor what every possessor of the skill has in common (*cf.* Thompson 2008). Rather, if Small is right, the appropriate level for theorizing about skill is that of the acquisition, development, and manifestation of a particular skill within an individual.

This brings us back to determination. Small insists that we interpret Locke's (1974) principle for type-individuating practical capacities as being relativized to the features of the constitution and background of individuals. Otherwise, we find ourselves mired in absurdity. Suppose, for *reductio*, that Locke's principle is fully general and applies to capacities irrespective of who has them. Consider two agents, one who can run a 5-minute-mile as well as a 6-minute-mile and one who can only run a 6-minute-mile. Applying Locke's principle, the first agent manifests a single capacity because the very same features of her constitution and background that standardly bring about her running a 6-minute-mile are those that standardly bring about her running a 5-minute-mile. Both are *determinates* of her capacity to run. If so, then this is not a capacity that the second agent has, since, by hypothesis, he cannot run a 5-minute-mile. But it seems that by Locke's principle the second agent *does* have the same capacity as the first. After all, the features of the second agent's constitution and background that standardly bring about his running a 6-minute-mile are (at least very likely to be) the same in kind as those that standardly bring about the first agent's running a 6-minute-mile. At the very least, given our initial assumption, it would be strange to deny that both agents have the very same capacity, namely, the capacity to run a 6-minute-mile. In which case, given unrestricted application, Locke's principle would suggest both that the capacity to run a 5-minute-mile is and is not the same as the capacity to run a 6-minute-mile. To avoid this

absurdity, then, we ought to relativize the principle to the constitution and background (and, thereby, capacities) of particular individuals. On this way of thinking, the two agents have the capacity to run a 6-minute-mile in a way that is derivative of their each having a capacity whose profile is sufficiently similar to that of the other individual's. They are participants in the same social praxis, and this is exhibited in the resemblance of "shape" across those agents' capacities.

The evolution of an agent's skills into more determinate capacities is an essential part of the life cycle of a skill. Consider a proficient pianist. She has reached the stage where she has an ability to play. Moreover, she might be able to play well for her level of training. But in order to avoid plateauing, that is, to better reach mastery, the proficient pianist must elaborate her skill (Dreyfus 2006). She must acquire skills, say, for playing certain progressions of especially difficult pieces as well as master *rallentando*, trills, arpeggios, and so on. She must likewise acquire discursive skills for critically assessing her practice and performance as well as to reason well about the domain. She must be able to see where she can improve her play and must be able to develop plans and actionable strategies for maximizing her chances of improvement. Importantly, the collection of further determinations of her skill as well as how exactly she comes to critically assess and reason about her practice and performance will in the end be specific to her. Indeed, part of the maturation of a skill into the achievement of expertise is the evolution of what Small calls its "specific shape *in concreto*" (2017: 257). That is, an agent's skill comes to have a profile distinctive of that disposition and specific to that agent. This profile changes over time, becoming increasingly determinate through the agent's continued practice and performance. Expertise brings style in its train, and this is at least partly a function of the agent's quite literally shaping her own capacity over time.

So much for particularity. Moving on, each moment in the *a priori* unity of the life cycle of an individual's skill is indebted to the existence of a corresponding social praxis. This is because the rationality of the skill extends to its acquisition and maintenance. Were the acquisition and maintenance of a skill not rational then it is unclear how their primary manifestation in performance could become rational. That is, it is unclear how expert performance could be expressive of what I will call the agent's "cogent practical cognition" (more on this in Sections 1.3.3-1.3.4) in a way that is flexible and situationally appropriate relative to the domain of her skill. In that case, it is difficult to see how skills would differ from the manifestation of vital dispositions where the response turns out by dint of luck to be appropriate.

To get a sense of what I mean, suppose a fledgling actress with stage-fright is playing Fantine in *Les Misérables* and undergoes an automatic freeze response when, in the play, Fantine is arrested by Constable

Javert. Despite appearances, such a response would not count as a portrayal of Fantine at all but merely a failure on the part of the actress to perform appropriately. By contrast, consider a seasoned thespian's portrayal of Fantine's fright at being arrested. This latter's performance is precise and sensitive to the context of performance. Even supposing that the thespian's performance is qualitatively identical to the fledgling actress's freeze response, what makes the latter's an appropriate portrayal of Fantine is that it is a non-accidentally successful expression of her cogent practical cognition of how to play the character. Such expression is itself the realization of the relevant intentional action concepts and is reflective of the rationality involved in the acquisition and maintenance of the thespian's skill. The point for now is that if neither her training nor practice were rational, her portrayal of Fantine would really be no different from the fledgling actress's freeze response.[19] That the thespian's performance is distinct from the freeze response is in large part due to the thespian's being inculcated into and becoming a practitioner of the art of theatrical performance.

In fact, it is thanks to the existence of social praxes that instances of incorporation, subsumption, determination, and integration throughout the course of the life cycle of an agent's skill involve the acquisition of the relevant discursive skills in the first place. That is, the specific shape of an agent's skill *in concreto*—including those habits and skills that come to constitute parts of that shape—depends on the existence of a corresponding social praxis. This dependence of shape on corresponding social praxes exists despite the former's particularity. Again, according to Thompson (2008), social praxes unify the actions that are their exemplifications across individual performances and practitioners, act as standards for the normative assessment of those actions, and act as *explanans* for those actions as internal to the relevant praxis (199). These three features of praxes are what underwrite a skilled agent's coming to be sensitive to the standards of the domain of her skill as well as her ability to reason about how to do well or improve in that domain. These features of social praxes are likewise ultimately what allow skilled agents to develop their own standards, principles, methods, techniques, heuristics, and so on as additions to the edifice of the relevant praxes. Most importantly, it is these features of social praxes that ground the rationality of the acquisition and maintenance of an agent's skill.

If, as Small and Thompson have it, the acquisition and maintenance of a skill are rational then they must be articulable, explicable, and assessable by reasons. Such articulation, explanation, and assessment are in principle shareable. An agent can introduce these reasons to others such that she can unify, assess, and explain

[19] Ryle (1949) makes a similar argument concerning the well-timed tumble of a skilled clown: were his training not rational then he would really be no different from the clumsy man whose actions the clown's performance perfectly mirrors (21-24).

those others' performance, practice, and teaching (and *vice versa*) as internal to the resulting social praxis. The possibility of skill thus depends on the possibility of a corresponding social praxis. Indeed, Small contends that a skill's acquisition and maintenance must be rational and from this that "it follows that there is a skill only where there is a [social praxis]" (103). He goes on:

> [that acquisition and maintenance are rational] means more than that no one can have the skill for doing *A* unless there are, at least potentially, other *A*-ers; it means that the [social praxis] is self-sustaining, that it is the cause of its bearers [...] the skill must *keep itself alive,* through the rational transmission of itself (the form of its activity)—that is, through teaching (103-104; original emphasis, variables replaced for consistency).

Part of what makes a skill rational, then, is that its transmission is also rational and that this is due to the existence of a corresponding social praxis. This fact bears directly on the intelligence of these dispositions. What's more, the same is true of habits: what makes an intelligent tendency intelligent is that its transmission is also intelligent. Good habits thus likewise depend on there being a corresponding social praxis supporting their inculcation. It is to this that I now turn.[20]

1.2.4 Rote Routine versus Intelligent Tendency, Mere Competence versus Expertise

Small (2021) holds that a distinguishing feature of both intelligent tendencies and skills is the possibility of rational transmission and, thus, the possibility of their being intelligent. Small draws a fourfold distinction with respect to the intelligence of these dispositions (Figure 1.3). On the one hand, volitional dispositions range from unintelligent rote routines to intelligent tendencies. On the other hand, rational dispositions range from mere competences to expertise concerning some domain. Incorporation, subsumption, and integration are relations that habits and skills stand in to each other due to their share in intelligence. And determination of the shape *in*

[20] One might worry that requiring that skills be backed by social praxes precludes nonhuman animals from having skills. Yet, it seems like many such animals behave skillfully in that they improve their capacities for action on a number of dimensions through practice, imitation, and so on and seem to teach their young some of these capacities, e.g., for hunting. In response, the powers framework can allow for a continuity between the poles of arational volitional dispositions in nonhuman animals and full-blown rationally articulable skill or expertise in humans. Nonhuman animals that communicate, form social bonds, and impart abilities to their young through demonstration and the like are transmitting capacities for action to their young. And animals who improve their abilities through experience are making those abilities more intelligent. Such transmission and learning can be indicative of capacities that we might well call "animal skills." The focus on rational transmission of skills and their dependence on social praxes is meant to highlight that such transmission and learning in us is transformed by rationality. The difference rationality makes in some way depends on a unity of the agent and praxis that is only made possible by discursive representation (see Section 1.3.1*ff*.). And while that difference is, to my mind, extremely important for understanding mental action and the human mind, it is not a difference that precludes evolutionary continuity between us and nonhuman animals. So, "animal skills" might well refer to evolutionary precursors or sub-capacities involved in human skills or in the rational transmission thereof. And it is a perfectly legitimate use of "skill" if so. Thanks to Colin Allen for pushing me to clarify this.

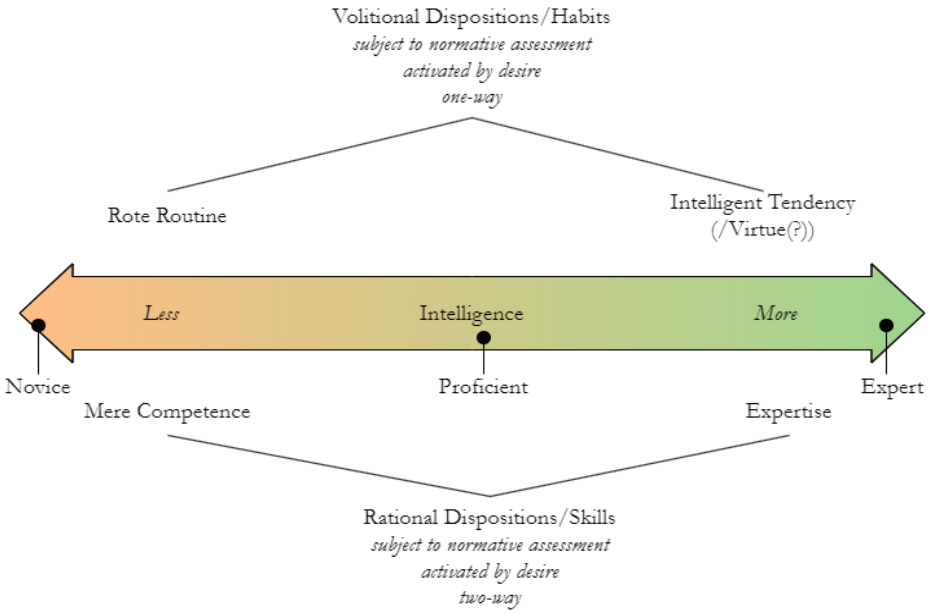

Figure 1.3 The Intelligence of Skills and Habits

A mixed diagram depicting the shared dimension of intelligence across which both volitional and rational dispositions range. To connect this diagram to Figures 1.1 and 1.2, I have included the stages of the life cycle of a skill, placing them along the spectrum of intelligence appropriate to the degree of intelligence required for achieving novicehood, proficiency, and expertise.

concreto of a skill is a function of its intelligence. The development of intelligent tendencies from rote routines and the development of expertise from mere competences, then, are interdependent to the extent that they involve instances of these four relations. Moreover, both types of disposition, if they are to have a share in intelligence, presuppose the existence of a corresponding social praxis. In this section, I will speak of the acquisition and development of rote routines and intelligent tendencies as well as that of mere competences and expertise. With the exception of Section 1.4.3, I will afterwards switch back to speaking just of habits and skills.

Starting with skills, the life cycle gives us a good model for how these dispositions become increasingly intelligent over time. Each skill starts its existence at acquisition as a mere competence and through instruction, guided and unguided practice, and performance matures into expertise. The proficient pianist has a mere competence to play but, as yet, cannot play *well*. Her mere competence is not sufficiently flexible or fine-grained: she might only be able to play a small selection of pieces, will not know how to perform more subtle and nuanced movements, and will not be able to easily integrate her mere competence for playing piano into expertise for playing harpsichord. She likewise lacks discursive expertise for properly explaining, assessing, or reasoning about her own performances or that of others. And she is not in a position to teach someone else how to play (beyond her level of competence) even if she has exemplary pedagogical expertise.

All of this lack stems from the fact that the intelligence of her mere competence is not yet perfected. The perfection of the intelligence of expertise consists in the agent's having comprehensive actionable procedural knowledge-how to engage in the relevant skill. Such knowledge, when it is connected to the ends of the corresponding social praxis, is the productive knowledge that the agent's expertise consists in and that she gains by means of training as well as both guided and unguided practice (Small 2021: 239-247). That is, such perfected knowledge consists in a comprehensively articulated network of intentional action concepts ready to be realized in skillful play or applied discursively in critiquing, assessing, or explaining one's own or others' performances. The master pianist plays skillfully, can immediately recognize a masterful performance, is well positioned to instruct others, and can purposefully play poorly relative to any number of dimensions of assessment or for ends contrary to the good of the domain.[21]

Turning to habits, Small (2021), following Ryle (1949), distinguishes between what I am calling "rote routines" and intelligent tendencies. Small does not spell out a life cycle for habits. However, I surmise that their intelligence depends on the intelligence of whatever mere competences or expertise are required for their exercise. This dependence is due to the fact that the perfection of the intelligence of intelligent tendencies consists in the agent's knowing *to* (or not to) perform actions exemplary of the relevant disposition. Focusing on the virtues, Small claims that "the generous person *knows to* put others first and the just person *knows not to* take more than her fair share" (2021: 247; original emphasis). Each agent's developing these "intelligent dispositions of the will [...] or power of choice" presupposes her having expertise for doing what she wills or chooses to do in exemplifying the relevant intelligent tendency (2021: 247). If so, then the shape *in concreto* of the corresponding skill must be sufficiently determinate such that the agent can exemplify the relevant intelligent tendency when and where appropriate. That is, the agent must have expertise for performing some action for the intelligent tendency to perform it to develop. This holds just as much for the development of a virtue like generosity as it

[21] Note that, as stated, the explanation for the difference between the (merely) proficient and the expert is circular: the proficient agent lacks a fully realized skill because she lacks perfected knowledge-how and the perfected knowledge-how of the expert just is his fully realized skill. But this circularity is not vicious. Indeed, it's of the sort that Aristotle is concerned with in explaining the development of the virtues. The same sort of solution to his quandary should thus work here, namely, the temporal and interpersonal dynamics of education dispense the air of vicious circularity. The novice needs guidance in the form of the perfected knowledge-how of some expert teaching him to begin his journey towards mastery (Small 2014). Moreover, the opportunity for guidance continues well past novicehood and proficiency into expertise. Experts regularly guide each other in sharing novel principles, methods, techniques, heuristics, and so on. They can likewise guide themselves by devising and applying such novel principles and the like in their own practice and performance. The point is that the development of perfected knowledge-how in one individual (usually) stems from the already perfected knowledge-how of at least one other individual. This is the analogue of the process of reproduction in organisms.

does for the development of a rote routine like tying your shoes or taking your normal commute. And it is this dependence of rote routine and intelligent tendency on mere competence and expertise that explains why, as was claimed in Section 1.2.1, skills comprise capacities for intentional actions, whereas habits do not. Rather, the latter are dispositions for action in virtue of being dispositions for immediately realizing some knowledge-how constitutive of a corresponding rational disposition. Incorporation of a mere competence or bit of expertise by a rote routine or intelligent tendency, then, is an emergent relation that necessarily coincides with the latter's formation. I will have more to say about this as it relates to so-called "basic action" in Section 1.4.2.

The point to draw from the foregoing is that the natural development of habits is itself part of the life cycle of the relevant skills. For example, the rote routine of driving one's normal route home is the product of over practicing driving that route. Such a habit is a common, relatively unintelligent volitional disposition among otherwise (merely) proficient drivers. Yet, incorporation occurs not just with rote routine but with increasingly intelligent tendencies as well. The soldier's act of bravery incorporates his expertise in combat. The master pianist's intelligent tendency to practice daily incorporates his expertise for playing. Focusing on the latter, the pianist's intelligent tendency to practice daily is intelligent at least in part because, in practicing, he is applying his perfected procedural knowledge-how to and in what he is doing. He looks and listens for mistakes, identifies passages that are difficult for him, makes sure that he is not getting rusty or too comfortable with certain movements, and so on. It is because i) his intelligent tendency to practice incorporates his expertise for play, ii) this expertise itself consists in his perfected procedural knowledge-how to play, and iii) his practicing maintains his expertise by providing opportunities for reappraisal and reassessment of his behavior that his intelligent tendency to practice is itself intelligence.

The transitions from novicehood to proficiency and proficiency to mastery depend not just on incorporation but on subsumption, determination, and integration as well. Starting with subsumption, let us modify our earlier example of the fledgling pianist who has a rote routine to assume an appropriate posture when sitting down to play. Suppose that she has this rote routine but no corresponding habit when she sits down to type at her computer. Still, acquiring this rote routine and developing it into an intelligent tendency (along with several others) are integral to her acquiring the mere competence for playing and, thereby, achieving proficiency. She will not acquire a mere competence for playing if positioning herself, coordinating her hands, exhibiting finger independence, and so on are on every occasion things she must actively synthesize in

deliberation prior to or as she does them.[22] The need to develop a battery of rote routines or intelligent tendencies that come to partially underwrite acquisition is a way in which the intelligence of mere competence and expertise is dependent on the intelligence of rote routine and intelligent tendency. The fledgling pianist in the process of acquiring the mere competence for playing must know *to* sit thus-and-so, to move her hands thus-and-so, to move her fingers thus-and-so, and so on at least well enough to engage in some amount of play as part of guided and unguided practice and performance. When the pianist achieves proficiency, her increasingly intelligent tendencies to do these things are subsumed into the shape *in concreto* of her developing skill for playing.

Yet, when subsumption occurs, the subsumed rote routines or intelligent tendencies do not themselves become determinates of the subsuming mere competence or expertise. The relation of subsumption differs from that of determination because the types of dispositions constituting the *relata* of the former have distinct forms. The proficient pianist's rote routine or intelligent tendency for placing her hands thus-and-so before play is a volitional disposition, whereas her mere competence or expertise for positioning her hands thus-and-so is a rational disposition. The former are one-way while the latter are two-way. To see what sort of difference this difference in form makes, consider a proficient pianist instructing another on how *not* to position their hands by demonstration. Doing so requires successfully resisting manifesting the rote routine or intelligent tendency for positioning her hands correctly. She nonetheless exercises her mere competence or expertise for positioning her hands correctly in demonstrating how *not* to do so. In exercising the mere competence or expertise, she blocks manifestation of the incorporating rote routine or intelligent tendency.

By contrast, consider a master pianist who has a determinate bit of expertise, say, for playing in a pop-rock style as part of the shape *in concreto* of her skill for playing but nonetheless does not (yet) have an incorporating rote routine or intelligent tendency to play in that style. In that case, teaching another to play in a classical style does not involve any conflict between the exercise of the skill for playing and the lack of exercise of her determinate expertise for playing in the pop-rock style. She does not have to impede the exercise of this more determinate capacity. Rather, in such a case, the master pianist is exercising her more general skill for playing in instructing another along one of its "tracks," that is, the one which manifests in her playing in the classical style, and not another, that is, the one which manifests in her playing in the pop-rock style.

[22] Paraphrasing Wu (2013), what would be unbearable is not lots of automaticity in action but, rather, its complete absence.

Upon acquiring the mere competence for playing, the mere competences or expertise of the proficient pianist for positioning her hands correctly and the like become determinates of the shape *in concreto* of her newly acquired competence. That is, these other capacities give her newly formed mere competence its initial determinate shape *in concreto*. They are integrated into the new competence and the new competence is determined by them. By contrast, the rote routine or intelligent tendency for positioning her hands correctly cannot be a determination of her mere competence for playing. It is nonetheless part of the initial shape *in concreto* of that competence in virtue of being subsumed. Such subsumption is, in turn, due to the fact that her rote routine or intelligent tendency for positioning her hands correctly incorporates the corresponding mere competence or expertise for such positioning, possession of which partially underwrites acquisition of the mere competence for playing. Incorporation and subsumption, then, are mutually reinforcing relations between habits and skills that reveal their interdependence with respect to the life cycle of a skill and allow for a more fine-grained gloss of that cycle.

Finally, integration of habits and skills as part of the acquisition of a new skill likewise reveals the interdependence of skills both on other skills as well as on habits. For it is often the case that an agent begins gestation of a new skill already equipped with some others. Most of the time, we are in the position of our original fledgling pianist, possessing organized complexes of habits and skills ready to be given ampliative application in a new domain. In fact, it is only when a human being first becomes a rational agent that gestation of her first skills occurs in the absence of her possessing any prior habits or skills to be given such ampliative application. And it is only when she first becomes a rational agent that acquisition of those initial skills involves only incorporation, subsumption, and determination. Integration, then, is nearly as ubiquitous as the other three relations and its becoming a possibility for a person is a mark of her transformation into a rational agent. I will have more to say about this in the following sections and in chapter 2.[23]

[23] Transformative theories of rationality hold that the transition from nonhuman animal cognition to human mindedness constitutes a break wherein the perceptual, mental, and agential faculties of human beings are shot through with rationality (for instance, see Boyle and Lavin 2010; Boyle 2016, 2022). By contrast, additive theories of rationality hold that the transition from nonhuman animal cognition to human mindedness is a steady progression that involves the addition of capacities that animals lack but which, in us, do not fundamentally change the capacities that we share with animals. Despite claiming that there is a break between nonhuman animal cognition and human mindedness, transformative theories can nonetheless acknowledge a transition and, thus, a continuity, between the two. Springle's account of intentionality can provide some of the materials for such an acknowledgment, although I do not endeavor to show that in this dissertation (see Section 1.3.1).

For now, let us reconsider the life cycle of a skill as exhibited by someone already in possession of rational agency. Such an agent starts with at least some rote routines and mere competences outside of the relevant domain. She is guided (usually) by a sufficiently skilled practitioner to give those capacities ampliative application in the acquisition of some mere competences in that domain and, with them, some plausibly rote routines that incorporate those new competences. In novicehood, the development of those rote routines in guided and unguided practice as well as performance scaffolds the acquisition of yet further mere competences that subsume those rote routines. These further competences include relatively rudimentary discursive ones for assessing practice and performance relative to the domain. Manifestation of these rudimentary discursive competences in practice is directed onto the subsumed rote routines and incorporated mere competences and allows the agent to identify and parse her old routines and assemble or "chunk" them into more intelligent tendencies of her now emerging skill.[24] Manifestation of these more intelligent tendencies and the mere competences they incorporate in practice and performance, in turn, scaffolds the acquisition of yet more flexible and fine-grained competences that subsume those more intelligent tendencies and further determine the shape *in concreto* of their incorporated competences. These mutually reinforcing sub-cycles of incorporation, subsumption, and determination proceed along the spectrum of intelligence from rote routine and mere competence to intelligent tendency and expertise. Yet, such reinforcing sub-cycles are present even at the developmental stage prior to proficiency.

At a certain point, the novice achieves proficiency and what was before a disunified complex of competences and tendencies are finally integrated into a competence that is definitional of the skill. In other words, at proficiency, the agent acquires the skill that the corresponding social praxis has the function of inculcating. The shape *in concreto* of this new skill will consist in, on the one hand, the sub-skills she thus far acquired as determinates and, on the other hand, the habits she has thus far developed as subsumed constituents. She can now go on to determine this shape *in concreto* through the acquisition of yet further determinate sub-skills and the subsumption of yet further habits.

Moreover, the now proficient agent will be able to turn her increasingly articulate discursive skills onto discursive habits that she developed as part of achieving proficiency. That is, the way that she tends to think about her practice and performance will now become an object of her thinking and, thus, subject to critique, assessment, and explanation. She may find that her teachers or their techniques were not perfect models of the

[24] At the very least, she should be in a position to identify bald mistakes and to take steps to correct (for) them.

social praxis of which she is now herself a practitioner. Or she may find that the social praxis is in need of something new. Or she may find that, although her teachers were good and she sees no reason to push the boundaries of the social praxis, she would do well to think of that praxis in a new light. At this stage, the agent is reasoning about the ends of the social praxis into which she has been inculcated. In any of these cases of reasoning about ends, how she is able to think about her skill is likely to shift and the way that she tends to think about the skill is likely to incorporate that shift. She might then become aware of new possibilities for the further development of her skill, including the discovery of new principles, methods, techniques, heuristics, and so on. Her adoption of these principles and the like will, in turn, result in her acquiring mere competences or expertise and rote routines or intelligent tendencies reflective of the development of a unique style. Such stylistic capacities will subsume both new and old habits and, with practice, the resulting stylistic habits will incorporate a range of new and old skills. By now, the agent is achieving mastery and becoming an expert. In developing her own style, she has added to the edifice of the social praxis into which she has been inculcated. And, assuming she has a modicum of pedagogical skill, she could well teach someone else what she knows.

Note that this more fine-grained gloss on the life cycle of a skill keeps the development of skill primary. Note too that it is sparse concerning exactly what skills and habits the agent develops along the way to becoming an expert. Small is correct to prioritize skill. After all, a habit to do something presupposes the capacity to do it, whereas the capacity to do something does not presuppose a corresponding habit. Yet, if this new gloss on the life cycle of a skill is right then acquisition of any skill presupposes at the very least possession of a battery of rote routines or intelligent tendencies and their incorporated competences or bits of expertise. Indeed, even Ryle acknowledges that "naturally skills contain habits" (1946/2009: 234). A wrinkle in the interdependence of the intelligence of both habits and skills that this new gloss makes explicit is the particularity of the latter's developmental trajectory. While the existence of a social praxis secures unity and standards across practitioners, as this new gloss shows, each individual's skill has a unique shape *in concreto* due to the history of its development. In particular, the uniqueness of determinate shape *in concreto* is ultimately due to the agent-specific history of the sub-cycles of incorporation, subsumption, determination, and integration involved in that agent's training, guided and unguided practice, and performance.

An upshot of this more fine-grained gloss, then, is that the intelligence of subsumed rote routines or intelligent tendencies and of determinate competences or bits of expertise vary between practitioners. Of course, there are some habits that are fundamental in the sense that they are shared across (nearly) all practitioners. These constitute the basic actions of the relevant domain. For instance, in the case of piano playing, positioning

one's hands correctly is a habit that is (nearly) universal across practitioners. And so positioning one's hands is something practitioners can do without having to do anything else. But these fundamental habits are not usually all that interesting or informative of the nature of the skill containing them. Rather, they are merely requirements on acquiring the skill in the first place. By contrast, as the new gloss brings out, mastery brings with it discernible uniqueness of shape *in concreto* through the interplay of incorporation, subsumption, determination, and integration.

Another, related upshot of the more fine-grained gloss on the life cycle of a skill is that the tracks of the relevant skill will likewise vary between practitioners. Following Ryle (1949) and Small (2021), skills are multitrack dispositions in virtue of being complexes of other mere competences or bits of expertise and rote routines or intelligent tendencies. These complexes form networks parts of which the agent possessing the skill can differentially manifest. Given the history of each expert's development, some experts might well lack determinate sub-capacities that others have. In that case, the tracks of a skill differ between individual practitioners who nevertheless possess the same general capacity. Of course, there will be a set of minimal tracks that (nearly) every fully realized practitioner is able to manifest in manifesting the skill. But, again, these tracks are usually not all that interesting or informative of the nature of the skill. Rather, they are merely the early, enduring infrastructure upon which the rest of the skill is built. By contrast, as the new gloss brings out, what is of interest in understanding the nature of skill is making sense of the development of the whole (super)structure and the agent's differential activation of parts thereof.

Finally, a third upshot of this new gloss on the life cycle of a skill is that the priority of skills over habits is specific to incorporation. Skills are not in general prior to habits or *vice versa*. An agent seeking to acquire a new skill does not start out either with only a set of mere competences or with only a set of rote routines. Rather, the mutual interdependence of habits and skills speaks against thinking of the former as in general less intelligent. Indeed, bits of expertise are often incorporated into corresponding intelligent tendencies as part of achieving mastery. Bits of expertise and fully intelligent tendencies are thus on a par when it comes to their share in intelligence.

So far, my focus has been squarely on the metaphysics of the powers framework. Skills are rational dispositions and habits are volitional dispositions. Rational dispositions are complexes of other rational dispositions and volitional dispositions. The rational dispositions composing a skill are determinations of that skill, whereas the volitional dispositions composing a skill are subsumed constituents of it. Those subsumed volitional dispositions incorporate some of the determinations of the skill. And the determinate rational

dispositions further subsume volitional dispositions that are ultimately subsumed by the skill itself. In most cases of skill acquisition, at acquisition, the new skill integrates some of the volitional and rational dispositions of some other skill(s) that the agent already possesses. Except in Section 1.4.3, in what follows, I will go back to speaking of habits and skills rather than rote routines or intelligent tendencies, on the one hand, and mere competences and expertise on the other.

I have thus far simply taken for granted that the intelligence of skills is embodied in their constitution. I have likewise taken for granted that the intelligence of intelligent tendencies is embodied in their constitution. Finally, I have taken for granted that this intelligence consists in the intentionality of the component dispositions of skills and intelligent tendencies and in the rational relations in which those component dispositions stand. All three assumptions concern the *practical* intelligence that at least partly distinguishes the powers framework from its causalist rivals. Understanding practical intelligence is key to understanding the structure that exists within and across skills and habits. Moreover, understanding such intelligence is key to understanding how intentional action is at once the expression of the fundamental building block of skill and habit and, yet, dependent on skill and habit for its very possibility. To account for the intelligence of intelligent agency as understood by the powers framework, then, I must provide an explication of practical intelligence. It is to this that I now turn.

1.3 The Nature and Intentionality of Practical Intelligence

Earlier (Section 1.2.4) I said that expertise in some domain is perfected procedural knowledge-how (not) to perform actions constituting exercises of the relevant skill and that a fully realized intelligent tendency is perfected knowledge (not) to perform actions constituting exercises of the relevant habit. These dispositions are kinds of knowledge in virtue of being relatively enduring aspects of an agent at the personal-level that bear both systematic normative, rational relations to their objects and systematic practical-inferential relations to each other. That is, they are kinds of knowledge because, in effect, they represent their objects in a particular way, are non-accidentally factive, and cohere with each other in a practical, rational way. That expertise and intelligent tendencies are perfected forms of knowledge is indicated by the fact that the success of an agent's exercising such expertise or such a tendency tends to be a matter of course, whereas failure requires appealing to something circumstantial. For instance, when a master pianist makes a mistake in playing a piece, it makes most sense to look for an explanation in something other than her expertise. She got distracted, was not feeling well, (if we're feeling inventive) was drugged, or… One does not look for any such explanation of her success. Rather, her

success is attributable to her possessing the relevant expertise itself. By contrast, one also does not look for such an explanation when a novice makes a mistake. Mistakes are to be expected from a novice and from a merely proficient agent. And this is precisely because their procedural knowledge is not yet perfected. They do not (yet) know how to do those things which exhibit mastery in the relevant domain.[25] Similarly, it makes most sense to look for an explanation in something other than the intelligent tendency and its exercise for what went wrong, say, if a genuinely generous person did something miserly. One does not look for any such explanation of her success. Rather, her success is attributable to her possessing the virtue of generosity right away. One does not look for such an explanation when a person lacking full possession of the virtue of generosity does something miserly. Failure to act generously is to be expected from the non-virtuous, as their knowledge (not) to perform certain actions is not yet perfected.[26]

Note that the objects of these kinds of knowledge are actions. I will argue in this section that the way that these kinds of knowledge represent their objects is practical rather than speculative. That is, unlike speculative knowledge, these kinds of knowledge are practical in that they represent their objects i) as within the agent's power, ii) as means to achieve some end or as ends to be achieved by the appropriate means, and iii) as to be carried out by the agent. That skills and habits are practical forms of knowledge and, thus, that their intelligence consists in a practical form means that to make sense of them we need an account of practical intelligence.

---

[25] The perfection of one's procedural knowledge of a domain through practice and performance explains the non-accidental factivity of the manifestations of that knowledge in action. Moreover, perfection is relative to the ability in question. A novice might have perfected procedural knowledge-how to position her hands correctly and mistakes in her doing this might warrant searching for an explanation of failure in the circumstances. The failure of someone lacking this sub-skill of the skill for playing piano to actually position his hands thus-and-so would not warrant such a search. Finally, acquisition and mastery of the relevant skill come to take the place of mastery of the relevant sub-skill in the order of explanation for non-accidental success and, thus, non-accidental factivity. That is, in the proficient and master pianists, the capacity to position one's hands correctly is incorporated into the corresponding habit to do so and that disposition is subsumed as part of the shape *in concreto* of the skill for playing. In that case, one appeals to the skill for playing and its shape *in concreto* in explaining the proficient and masterful pianists' non-accidentally successful and, thus, knowledgeable hand-positioning. And this is because mastery of the skill for playing is the ultimate *telos* of acquisition and mastery of the sub-skill for positioning one's hands correctly across the contexts of doing, practicing, and teaching. The procedural knowledge-how constitutive of the latter in any individual owes its existence to the possibility of the procedural knowledge-how constitutive of the former. The possibility of knowledge-how to play in turn depends for its existence on the relevant social praxis.

[26] Following Anscombe, we might say that the non-virtuous in general fail to be in full possession of what Aristotle called "practical truth" (1999). There is thus a sense in which any action that the agent conceives as falling under the heading of what is "befitting of a human as such" but which is not in point of fact an instance of "doing well" is not intentional because of a corresponding lack of practical knowledge: the agent does not exercise practical knowledge of (because she is not making true) her doing well. Thanks to Alison Springle for pushing me to clarify this point.

This brings us to the second thread running through the powers framework. Again, this thread originates in the work of Anscombe (1957/2000) and is concerned with the nature of intelligent agency writ large. According to philosophers working in this thread, intentional action is the paradigmatic expression of intelligent agency. And in paradigmatic instances of intentional action, namely, those performed self-consciously by neurally intact human beings "grown to the age of reason" under normal conditions, the agent is doing something knowingly where this means knowing *what* she is doing, *why* she is doing it, and *how* she is doing it (Anscombe 1957/2000: §4). Intentional action, the intentions with which actions are performed, and expressions of intention are Anscombe's tripartite target. Anscombe identifies the presence of an articulable calculative structure of means and ends as that which distinguishes the intelligence of intelligent agency, makes actions intentional under the relevant descriptions, provides the intention(s) with which an agent acts as ends calculatively connected to her means, and provides the content for the expression of intention. In central cases of intentional action, the agent is in possession of this structure in what some of Anscombe's followers calls her "knowledge in intention" such that she is in a position to say of herself, "I am doing *D* by doing *C*, doing *C* by doing *B*, and doing *B* by doing *A*," or, what amounts to the same, "I am doing *A* in order to do *B*, doing *B* in order to do *C*, and doing *C* in order to do *D*," where the variables (doing) *A*, *B*, *C*, and *D* take intentional action concepts as values.

Note that the series characterized by the by-relation gives voice to the agent's knowledge of *how* she is doing what she is doing. The application of the practical rational norm of means-end coherence depends on the possibility of this knowledge, which just is the agent's procedural knowledge-how. Note, too, that the series characterized by the in-order-to relation gives voice to the agent's knowledge of *why* she is doing what she is doing. The practical rational norm of responsiveness to what are actually reasons for performing the corresponding action depends on the possibility of this knowledge. It would be apt to describe this knowledge as the agent's knowledge of her reason(s) for action or, what amounts to the same, her knowledge of her practical reason(s). Such knowledge bottoms out in characterization of some end as desirable or as falling under some good. Such a characterization is appropriately responsive when the thing in question in fact falls under the relevant good.

Some action theorists hold that the calculative structure that Anscombe identifies is constitutive of intentional action and itself constitutes the agent's cogent practical cognition concerning what she is doing (for instance, Small 2012, Ford 2015, 2016). These philosophers differ on what exactly practical cognition ultimately consists in but agree that, whatever it is, such cognition distinguishes intelligent agency from non-

intelligent agency. The successful exercise of such cognition is the embodiment of practical rationality either in deliberation or intentional action. Deliberation applies intentional action concepts through the construction and use of what I call "discursive representations" that contain those concepts as their content. By contrast, intentional action realizes those concepts in making true the order that they set out as orders of means and ends that can be made true by the agent or anyone else situated as the agent is. That is, deliberation applies action concepts in synthesizing them into a discursive representation that it would be apt to call a plan. By contrast, intentional action instantiates the relevant order in the course of the relevant happening. This thread in the philosophical literature is concerned with explicating the nature of intelligent agency and its implications for the metaphysics of agency in general.[27]

In this section, drawing from Anscombe as well as some of her followers, I lay out what I take to be the distinguishing features of practical intelligence and that of its successful realization in practical knowledge. In the following section, I bring together the metaphysics discussed in Section 1.2 and the account of practical intelligence. It will turn out that skill and habit can be characterized in terms of intentional action concepts, that an agent's practical knowledge relative to a domain can be characterized as itself a sub-skill for correctly ordering the exercise of the relevant skills and habits within that domain in a way that can be rationally articulated into means and ends, and that the exercise of that practical knowledge just is the well-ordered manifestation of those skills and habits in intentional action (1957/2000: §§33-44, 48). However, to get there, I will be building up from the intentionality of intentional action concepts, to the "*generans*" and form of practical reason, to the intelligence of cogent practical cognizing, and, with such intelligence, the possession and exercise of practical knowledge. It is to doing this that I now turn.

1.3.1 The Intentionality of Intentional Action Concepts

According to the powers framework, practical knowledge is a domain-specific capacity to order doings within the relevant domain into means and ends sufficient to achieve, secure, or get whatever the agent aims to achieve, secure, or get. That is, on the powers framework, practical knowledge is itself a sub-skill that forms part of every

[27] Some action theorists deny that the agent must know what they are doing for their doing it to count as an intentional action under the relevant description. These philosophers treat as counterexamples to this claim cases like Davidson's carbon copier. In the original case, an agent makes 10 carbon copies without knowing that she is doing so (Davidson 1980/2001: 50, 82). Assuming she succeeds, she does what she intends, namely, she makes 10 carbon copies. But, surely, she cannot be said to know that this is what is happening as she is making them since the task is so difficult and the agent is liable to fail to press hard enough to mark the 10th copy. Or so these philosophers argue. I return to the case in Appendix D and to the controversy surrounding the so-called "knowledge thesis" or "cognitive condition" in Appendices C-E.

skill and habit the agent possesses and has a kind of domain-generality insofar as it extends as a skill for ordering doings from one domain to another. The exercise of practical knowledge is just the "supervising or doing" of those things of which she has that knowledge, which instantiated order supervised or done is an intentional action (Anscombe 1957/2000: §48). For practical knowledge to be contentful and, in the human case, for it to be rational, the capacity must stand to its objects in the appropriate way(s) and must be internally structured such that its parts bear the appropriate rational connections to each other. The powers framework contends that the contents of practical knowledge are what Small (2012, 2022) calls "intentional action concepts" or "doables" and which, following him, I will represent as "I/One can do *B* (by doing *A*)" and as "doing *B* (by doing *A*)." Which formulation is appropriate depends on whether the agent is performing an intentional action (I am doing); deliberating (I can do); describing, explaining, or assessing her own action (I am doing); describing, explaining, assessing, or supervising someone else's action (One is doing); or simply is in possession of the relevant concept(s) (I/One can do) (2012: 135). The specific content of an intentional action concept is the value of the variables "doing *B* (by doing *A*)" and refers to a means-end order that is within the agent's power to affect or which she is affecting in acting intentionally and that any agent situated as she is has the power to affect or is affecting in acting intentionally.

The question, then, is what such content consists in beyond this initial formulation. For if the powers framework holds that such content consists in internal truth-conditional representations that function to compose propositional directive attitudes then the framework is not a true alternative to its causalist rivals. On such a view of the content of intentional action concepts, an agent's behavior counts as an intentional action just in case she forms a propositional directive attitude applying the relevant intentional action concepts with the form "I am doing *B* (by doing *A*)" and that attitude bears the right productive and truth-conditional relations to the behavior, namely, that of (non-deviantly) causing or guiding it such that it makes true the content of the attitude. And this claim just is the root of at least contemporary causalism: the relevant attitudes are what causalists identify as belief-desire pairs, intentions, plans, goal-states, and so on. To distinguish itself from causalism, then, the powers framework needs an alternative conception of the content of intentional action concepts and, with it, an alternative conception of intention.

To my mind, there are four key moves required for the powers framework to distinguish itself from its causalist rivals when it comes to the intentionality of intentional action concepts. The first is to identify the content of intentional action concepts with the internally calculatively structured skills and incorporating habits themselves. The second is to identify the primary way in which such content is tokened with the activation of

the relevant part(s) of the relevant network constituting the internal structure of those dispositions. An important part of this move is holding that intentional actions are teleologically unified developing processes and that activation of the relevant intentional action concepts is just the initial stage of such processes. If this is right, then practical judgments consist in activation of systematically means-end ordered networks. The third key move is to identify as a secondary way in which such content is tokened with the formation of representations with truth-conditional contents that synthesize intentional action concepts together. If this is right, then articulation of practical judgments in thought or language are not identical to those judgments but, rather, are expressions of them. To remain consistent, the framework should then hold that the formation of such representations is itself the activation and subsequent exercise of specifically *discursive* parts of the relevant network. Such parts are what I am here identifying as deliberation and the discursive skills involved in expressing intentions as well as in describing, explaining, or assessing one's own or another's action. To do any of these things is to practically judge to express, say, in thought or language a distinct practical judgment. Finally, the fourth move is to hold that the rationality of the order realized in intentional action and synthesized in the formation of the relevant truth-conditional representations stems from the transformative effect of the activity of discursive sub-skills in shaping capacities for action into internally calculatively structured skills. The effect of such discursive activity just is the rational transmission of skill as part of its life cycle. The rest of this section and the next is, in part, the making of these four key moves.

Starting with identifying the content of intentional action concepts, I draw primarily from the work of Springle (2019, 2021), Springle and Buccella (2024), and Springle and Humphries (2021). Springle in particular seeks to provide a naturalistic, non-reductive, action-first account of mental representation primarily in nonhuman animals. According to her (as well as Small, Thompson, and others), intentional actions are teleologically unified developing processes. More specifically, according to Springle, intentional actions are processes for solving problems constituted by needs for flourishing, for instance, the need for nourishment or the need for edifying aesthetic experiences.[28] Springle calls this "the problem-solving account of intentional action." Like skills, then, intentional actions are themselves like living organisms in that they too develop in

[28] As the latter need suggests, the framework understands "need" broadly such that it extends beyond needs for mere survival or reproduction. In fact, it posits rational needs specific to the flourishing of human beings, for instance, aesthetic, intellectual, and certain sorts of social needs, as well as distinctive mechanisms for solving them such as methods of teaching their young and symbolic systems employed in communication and explicit problem-solving, i.e., reasoning. That said, after discussing the intentionality of intentional action concepts, I will return to using "desire" and "(practical) want" as synonyms for "need" except where clarity dictates.

characteristic stages or "moments" constituting an *a priori* teleological unity. Unlike skills, the moments constituting the *a priori* unity of intentional action are not doing, practicing, and teaching. Rather, they are activation, ongoing performance, and completion as the satisfaction of some need. In fact, each moment constituting the *a priori* unity of the life cycle of a skill itself consists in teleological *a priori* unities consisting of their own moments. Like skills, the teleological *a priori* unity of intentional action is an *a priori* unity because any intentional action is what it is, necessarily, by being an interaction whose occurrence functions to solve instances of a type of problem through the occurrence of three formally distinct moments: activation, ongoing performance, and satisfaction of a need.

The stage of intentional action with which Springle identifies the tokening of mental representations in general is the initial or gestational stage, that is, activation. On analogy with the life cycle of oaks, Springle labels this stage the "intentional action acorn" and elsewhere identifies such a stage with what she calls the "potentiation" of an intentional action. In human agents, such potentiation is nothing other than the activation of the relevant part(s) of the network constitutive of some skill or incorporating habit of hers. By positing intentional action acorns, Springle conceives of the intentionality of mental representation in general and the intentionality of the performance and completion of intentional actions as two sides of the same coin. According to her, this is because they constitute three stages in the life of an individual action. In the order of existence, neither intentional actions nor mental representations come first. They are simply phases of a single unified developing teleological process. However, Springle's "problem-solving framework is action-first in that it treats intentional actions as prior to [mental] representations in the order of understanding analysis, meaning that an analysis of [mental] representation must mention intentional action" (Springle and Buccella 2024: 10).

Drilling down a bit, according to Springle, representations constitutively function to provide what she calls "practical access" to targets:

> **Representation:** R is a representation just in case R provides (or is supposed to provide) an agent with practical access.
>
> **Practical Access:** An agent has practical access to a target just in case the agent is in a position to respond to the target, where she is so positioned just in case she is primed to produce one or more actions that are directed at that target. The agent is primed to produce such actions just in case her capacities for producing these actions have been activated so as to potentiate instances of such actions.

Representations are supposed to provide practical access to a target by positioning the agent to produce an appropriate action with respect to that target. Representations succeed if they "practically fit" the situation (including the target) at which they are directed, that is, if they are appropriate.

The fundamental normative standard for intentional action is *appropriateness*. According to Springle, an intentional action is appropriate just in case, were it performed to completion, it would non-accidentally solve some instance of the type of problem it functions to solve, where such problems are constituted by the agent's having some need in the relevant context. Appropriateness is not a matter of actually successful performance. For instance, eating an apple is an interaction generated by a subject that, when successful, satisfies a need for nourishment. Such an instance of eating is appropriate if the subject has that need on that occasion, there is an apple that the agent can eat, and if her closing the distance between her present situation and her eating it through the relevant interaction *would* satisfy that need. According to Springle, intentional actions such as eating are type-individuated by the problem that they function to solve, as these specify the types of interaction and targets that define the type of solutions that a type of intentional action functions to constitute.[29] So, eating is individuated as a type of intentional action that functions to solve the problem of needing nourishment.

The different means by which different subjects go about eating—the different ways they go about closing the relevant distance—constitute more determinate solutions to the same problem. Together with their relations of means-end coherence and their relations of coherence with the performance of other actions, these different means form *solution systems* that are appropriate for producing different targets for eating. Hunting, scavenging, grazing, shopping, going to a restaurant, ordering delivery and so on are sub-solutions constituting distinct mechanisms for non-accidentally satisfying the need for nourishment and together constitute a solution system for addressing that need. Solution systems are thus constellations of connected capacities for action that can be partly or fully activated by an appropriate need and target relative to an appropriate environmental context. The facts that are instances of the problem that an intentional action functions to solve (what Springle calls "need-facts") together with the facts that are instances of the type(s) of targets involved in solving the relevant type of problem (what she calls "target-facts") constitute a solution-type's domain of application. Appropriateness of an intentional action, then, is a matter of the partial or full activation of at least part of a solution system relative to a domain of application where the complete application of the solution and its sub-

[29] That said, a single action performance may function to solve more than one problem and thereby constitute two distinct solutions (the two-birds one-stone phenomenon).

solutions to that domain would, on that occasion, result in the satisfaction of the relevant need. Practical access, in turn, is a matter of the agent's being primed to produce appropriate intentional actions on the relevant occasions.

One last wrinkle as it relates to appropriateness is the relations of coherence between capacities for action other than means-end coherence. For Springle, such coherence consists in relations constitutive of what she calls "application accountability." Application accountability manifests in nonhuman animals in three ways: i) in corrections in the application of solution systems, that is, in corrections to ongoing performance (what Springle calls "vertical flexibility"), ii) in the acquisition or extinction of such systems according to whether they solve the relevant problems reliably enough, and iii) in the prioritization or suppression of the application of such systems as required for maintaining a practically coherent superstructure of those systems (what she calls "horizontal flexibility"; Figure 1.5B-D). These three types of manifestation of application accountability are three ways in which agents manifest responsiveness to practical reasons. In humans, (i)-(iii) manifest as a result of discursive activity, that is, the rational transmission of skills and explicit problem-solving or, what amounts to the same thing, (practical) reasoning. Such activity functions to improve specific solution systems as well as the superstructure in part by making reasons for action objects of representation, as reasons for *us* to do such-and-such. Rational transmission and explicit problem-solving, then, are self-directed intentional actions on this account. Moreover, such activity is that in virtue of which we represent successes and failures as such. Such representation, in turn, is a requirement on intelligent agency.

Returning to practical access, according to Springle, there are two key varieties of such access: unmediated practical access and mediated practical access. An agent has *unmediated* practical access to a target just in case she is positioned to produce an intentional action that is a direct response to that target, where a response is direct when its immediate target is none other than that target. An agent has *mediated* practical access to a target just in case the agent is positioned to produce an intentional action that is an indirect response to that target, where a response is indirect when its immediate target is something that takes the place of the original target. That is, an agent's practical access to a target is mediated when what she acts on directly thereby allows her to act on a target for which that thing "stands in" or "goes proxy" by virtue of her appropriately treating that thing as a stand-in or proxy. Corresponding to these two types of practical access are two species of representation:

> ***De Agendo* Representation:** R is a *de agendo* representation of a target, T, just in case R provides unmediated practical access to T.

> ***De Substituto* Representation:** R is a *de substituto* representation of T just in case R provides (or is supposed to provide) mediated practical access to T.

A *de agendo* ("about the action") representation is the activation of some capacity for unmediated actions. Such representations are potentiations of unmediated actions or "acorns" of such actions. When activated, such capacities constitute embodied non-descriptive modes of presentation of targets to nonhuman animals. That is, according to Springle, activation of capacities for unmediated action in nonhuman animals consists in perceptual activity whose intentionality, in turn, consists in "standing-for" the full application of the relevant solution system(s) to a target relative to some need that the animal might (but need not at that moment) have (Springle and Buccella 2024: 12). Such intentionality is coextensive with that of the corresponding activated capacities for action. As she puts it: "*de agendo* representations present agents with targets in terms of the potentiated actions that constitute them" (Springle and Buccella 2024: 12).

So, an animal's *de agendo* representation of an apple is as of something, e.g., eatable or to-be-eaten.[30] Its being so presented is a function of the at least partial activation of the animal's solution system(s) for eating in response to the obtaining of the target-fact of there being an apple and in light of the animal's coming into perceptual contact with the apple and (usually) its needing nourishment (Figure 1.4).[31] Drilling down on individual capacities for action, it is worth noting that, according to Springle, each capacity is composed of causal sub-capacities: a Need-Input sub-capacity, a Target-Input sub-capacity, and an Interaction-Output sub-capacity (Figure 1.4A). None of these sub-capacities possess intentionality on their own. Only the capacity of which those sub-capacities are parts possesses intentionality. And this is because it is only the activation and manifestation of a capacity for action that is (or fails to be) appropriate in the first place. Because actions are what are first and foremost subject to appropriateness conditions, the activity of the causal mechanisms underwriting them are subject to such conditions only *derivatively*. Putting the point another way, a capacity for action has intentionality in virtue of its having the function to solve instances of a type of problem by tokening interactions constitutive of the relevant action. Any such tokening is underwritten by the activity of an Interaction-Output sub-capacity when the relevant Need-Input and Target-Input sub-capacities are affected by the obtaining of a corresponding need and target, respectively (Figure 1.4D). The causal sub-capacities are unified by virtue of being components of an agential ability that functions to solve the relevant type of problem.

---

[30] Hereafter, I will just use the "-able" locutions for representations of things as *A*-able or as to-be-*A*-ed, except where clarity dictates.

[31] Springle's *de agendo* representations are meant to develop Millikan's (1995) notion of "pushmi-pullyu representations" and, thus, have affinities with what it is to perceive Gibsonsian affordances.

Yet, while activation of the whole capacity sufficient to potentiate the relevant action involves stimulation of parts of the animal by virtue of the impact of its environment (including its own internal states) on the Need-Input and Target-Input sub-capacities, such impact does not amount to intentionality independent of the whole capacity. Similarly, while manifestation of the whole capacity involves the activity of the Interaction-Output sub-capacity through which the animal impacts its environment (including its own internal states), such impact likewise does not *of itself* amount to intentionality independent of the whole capacity. Considered in the absence of the unity provided by the function *of the capacity* for action, the impact of the environment on the animal through the Need-Input and Target-Input sub-capacities and the impact of the animal on the environment through the Interaction-Output sub-capacity constitute mere tokenings of causal relations. Only when the activity of these sub-capacities together constitutes a purposeful interaction between animal and environment does the animal counts as directed at something in the world and, so, as representing it.[32]

There are at least four structural features of *de agendo* representations worth considering. First, *de agendo* representations can overlap and form series corresponding to a network of sub-solutions (Figure 1.5A). For instance, a terrestrial animal might have a solution system for eating that incorporates a sub-solution for closing the spatial distance between themselves and an apple on a nearby tree , say, by climbing the tree or shaking it. The *de agendo* representations involved here would present the apple as gettable and the tree as climbable or shakeable. The apple's being within reach would then afford the animal the opportunity to actually consume it. The series of *de agendo* representations here would fall under a solution system for eating. That is, such representations would *be* the at least partial activation of some part of the animal's solutions system for eating. Importantly, it is the possibility of overlap between solution systems with respect to shared sub-solutions that allows for consistency of content across performances. That is, Springle's framework can abstract a shared content across such instances and, thus, avoid an old problem that plagues behaviorism. Behaviorism's rejection of mental representation meant that it had to explain the ascription of mental kind terms in terms of dispositions to behave. But because many mental kind terms are associated with indefinitely many different behaviors, behaviorist were forced to claim that the ascription of a given mental kind term amounted to the ascription of an indefinitely long disjunction of possible behaviors that the subject in question

[32] Following a broadly Thompsonian approach, the unity of capacities for action derives from the unity of agents which, in turn, derive from the unity of the lifeform of which the agent is an instance. It is therefore the *telos* of the whole lifeform that grounds the propriety of the practical-inferential relations among an animal's solution systems as well as the domain of application of such systems. The lifeform thereby grounds the appropriateness of actions and, thus, the intentionality of *de agendo* representations by providing the underlying unity to which those actions, when appropriate, contribute.

may or may not perform. Indeed, this was an issue for Ryle (1949). By contrast, on Springle's view, not only is there a mental representation whose content can be extracted, but the same content can inform performances of distinct actions through application of the same overlapping sub-solutions. So, for instance, an animal that interacts with trees in very many different ways will, under normal conditions of coming into perceptual contact with a tree, be intentionally directed towards the tree. Regardless of what it goes on to do, it represents the tree as a target for some such action among its tree-involving solution systems. And this representational core is attributable to the animal across instances without having to list an indefinitely long disjunction of possible behaviors.

Second, among an animal's solution systems are those that are exploratory: appropriate application of these solution systems is a matter of generating more targets and opportunities for action either in the here and now or when it is appropriate. A new space affords exploration and such exploration might uncover further targets or opportunities for action, say, for eating, burrowing, nest-building, and so on which the animal might capitalize on immediately. Alternatively, the new space affords exploration which results in the modification of solution systems or in the development of novel solution systems to be applied later. These modified or novel capacities for action need not be long lived. The point is that an animal, say, a mouse, might explore a new space such that, later, when some food is put into that space, say, at the end of a maze, the mouse can then apply its modified or newly developed solution system(s) without having to re-explore the space. In fact, the modification and development of (novel) solution systems in light of changing need- and target-facts is arguably the rule rather than the exception for both humans and nonhuman animals.

Third, differential patterns of activation among competing solution systems can result in suppression or prioritization of some solution system over another. At least when such patterns of activation are application appropriate, the result is more nuanced *de agendo* representations. For example, a zebra at a watering hole might spot a jaguar nearby. It might go on to simply drink from the watering hole, run from the jaguar, or cautiously drink while repeatedly checking for the jaguar's location. The first response might well be the prioritization of the solution system for quenching thirst and the suppression of the solution system for fleeing immanent danger. The zebra might then *mis*represent its overall situation as safer than it actually is. Or it might be appropriately representing its overall situation as calling for consumption even at the risk of being attacked. The second response, in turn, might well be the prioritization of the solution system for fleeing immanent danger and the suppression of the solution system for quenching thirst. The zebra might then be *mis*representing its overall situation as more dangerous than it actually is. Or it might be appropriately representing the danger even

at the expense of a much-needed drink. The third response might well be partial activation of both solution systems. The parts of the fleeing solution system for maintaining vigilance are intermittently activated while the parts of the thirst-quenching solution system for drinking are suppressed and *vice versa*. The zebra might then be appropriately representing the situation as dangerous but not immediately life-threatening. In any case, the practical-inferential relations of (in)consistency of actions as well as of means and ends are reflected both in the structure of the solution systems themselves and in the intentional content of the *de agendo* representations constitutive of those systems' (partial) activation.

Finally, on Springle's account, as I understand it, nonhuman animals do not represent their own actions or their capacities for action. The activation of these capacities is what represents and such representation is transparent to the targets (and needs) with which the animal interacts in the complete application of corresponding solution systems. That is, on Springle's account, capacities for unmediated action do not contain sub-capacities for referring to the capacity itself. Activation of the capacity in light of the obtaining of a need might well involve interoception on the part of the animal and with it a *de agendo* representation of what is needed or the interaction that is supposed to get it. The animal might look for the desired object or for an opportunity to perform the relevant action. It might even feel a sense of urgency. But it does not refer to itself. At most, nonhuman animals represent the actions that it perform as targets in further *de agendo* representations. For instance, a beaver represents the gaps in the dam it is building fillable with wood, mud, stones, vegetation, etc. Such a representation further potentiates a search for these materials. Or the beaver represents the dam it has successfully built as a target for sheltering in. Activation of capacities to search for more building materials or to shelter in the damn just is the beaver's *de agendo* representing the dam as buildable

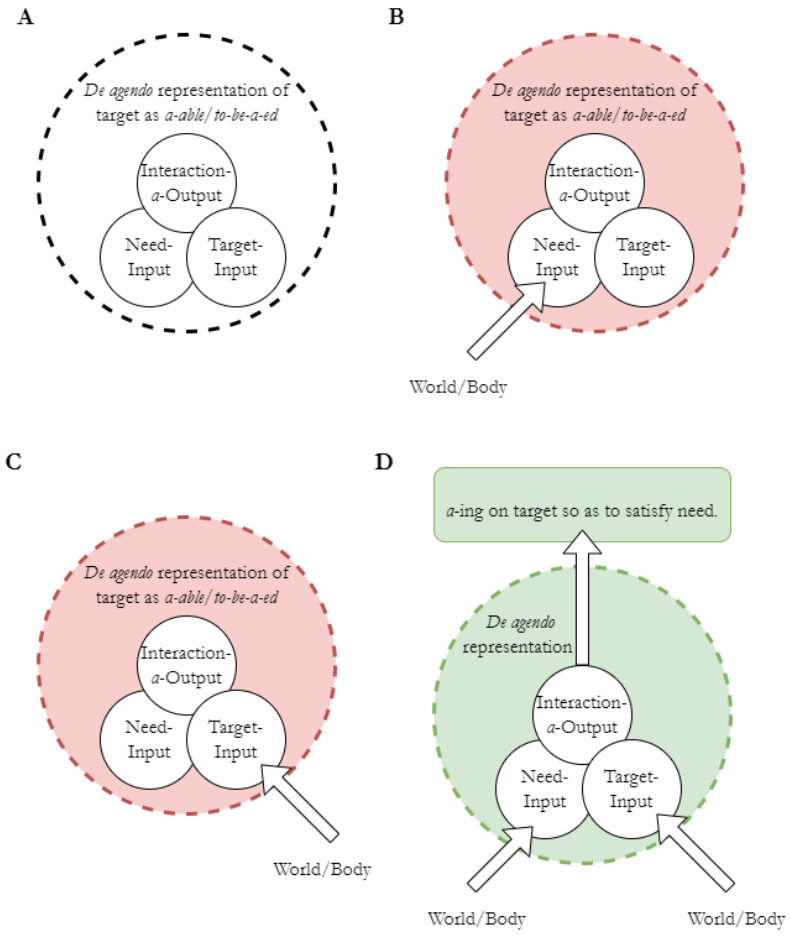

Figure 1.4 De Agendo Representations

A diagram depicting *de agendo* representations at various levels of activation. Red coloration represents activation in the absence of performance while green coloration represents activation sufficient for performance. **A.** Depicts an unactivated *de agendo* representation, or, what Springle calls a "*de agendo* concept" for generating a type of interaction (in this case, *α*-ing) with a type of target in response to a type of need. The Need-Input, Target-Input, and Interaction-*α*-Output circles represent causal sub-capacities of the capacity for *α*-ing. The doted circle enclosing them represents the unified capacity to *α*. **B.** Depicts a *de agendo* representation that is activated by an animal's registering a need. Such registration is analogous though not identical with the acquisition of a desire. **C.** Depicts a *de agendo* representation that is activated by an animal's coming into perceptual contact with a target. Such perceptually mediated activation is an animal's representation of that target as *α*-able. But the animal does not *α* for lack of the relevant need. **D.** Depicts a fully activated *de agendo* representation in light of registering the relevant type of need and coming into perceptual contact with an instance of the relevant type of target. Full activation represents the target as *α*-able and the representation will develop into the animal's *α*-ing the target so as to satisfy the need. Adapted from Springle and Humphries (2021).

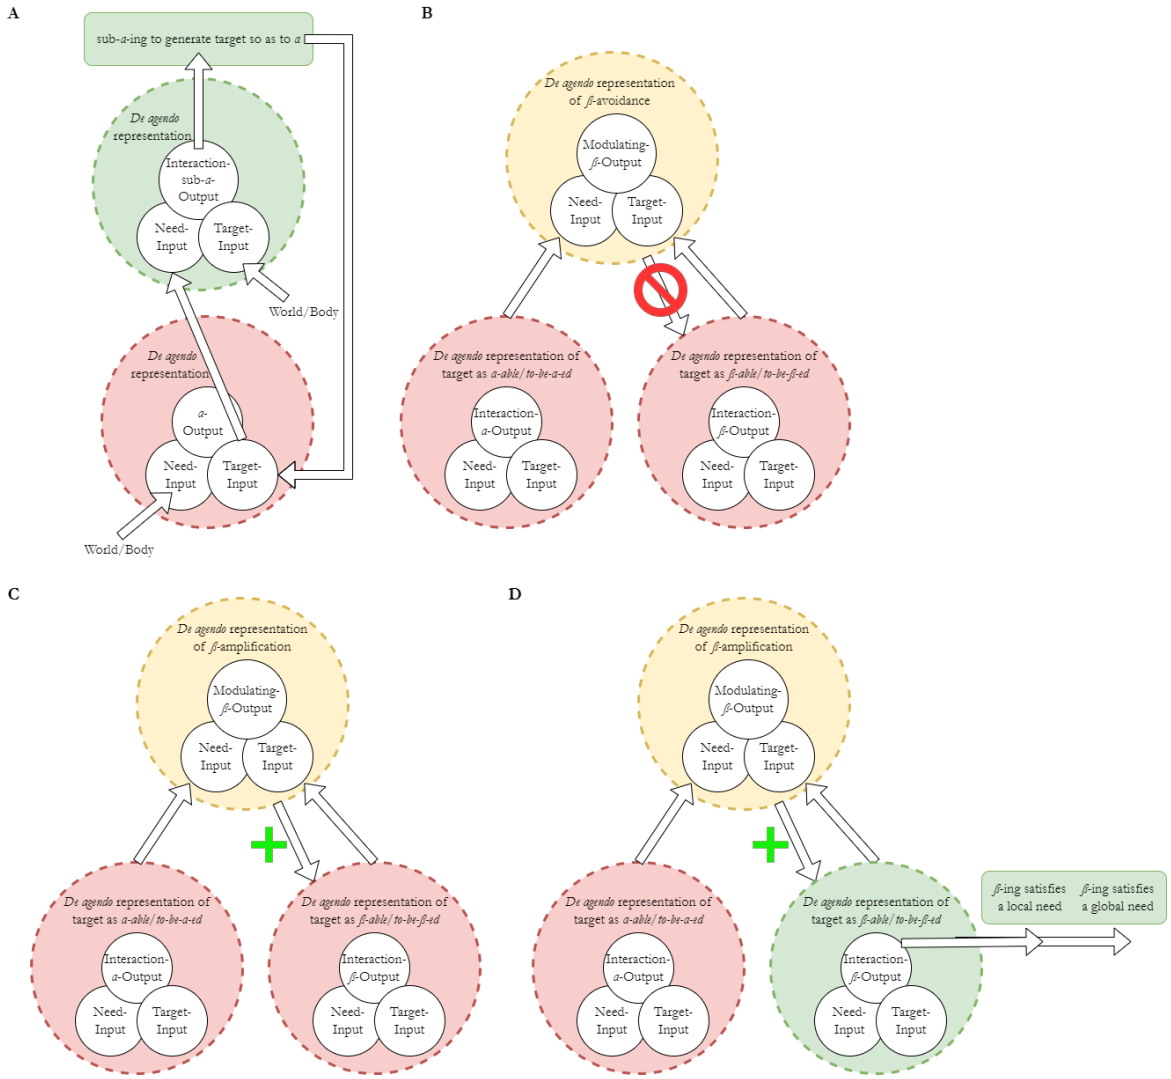

Figure 1.5 Animal Solution Systems

A diagram depicting various relations between *de agendo* representations constitutive of animal solution systems. Red coloration represents activation in the absence of performance. Yellow coloration represents activation sufficient to modulate the activation of other *de agendo* representations. Green coloration represents activation sufficient for performance. **A.** Depicts a *de agendo* representation whose activation by the animal's registering a need drives activation of a *de agendo* representation for discovering or generating the relevant target (what Springle calls a "terminal target"). The latter *de agendo* representation constitutes a sub-solution to the problem that *α*-ing on the discovered or generated target functions to solve. **B.** Depicts modulatory relations between three *de agendo* representations. In this case, an animal is presented with the opportunity to either *α* or *ß* but not both. In such a case, activation of both *de agendo* representations for *α*-ing and *ß*-ing drives activation of a modulatory *de agendo* representation (what Springle calls an "organizing *de agendo* representation") that, in this case, inhibits activation of the *de agendo* representation for *ß*-ing. In such a case, the animal will likely go on to *α*. **C.** Likewise depicts modulatory relations between three *de agendo* representations. The case is the same as that depicted in **B.** except that the modulatory *de agendo* representation drives further activation of the *de agendo* representation for *ß*-ing. In such a case, the animal will likely go on to *ß*. **D.** Depicts the function of such modulatory *de agendo* representations, namely, to promote overall flourishing by the animal's performing actions that not only satisfy some immediate need (what Springle calls a "local need") but that satisfy needs constitutive of that flourishing (what Springle calls "global needs"). Sometimes local and global needs overlap. Sometimes they do not. Adapted from Springle and Humphries (2021).

or shelterable-in. Yet, such activation itself is not the object of the animal's representing. Moreover, nonhuman animals do not represent the actions of conspecifics except as targets in *de agendo* representations. For instance, a sentinel meerkat's call is the target of another meerkat's *de agendo* representation wherein the former's call is represented as respondable-to. Which actions are afforded as a function of responding to the call will depend on a number of factors including how often the sentinel plays that role, the nature of the call, and so on.

Importantly, on Springle's account, these *de agendo* representations do not provide mediated practical access to the targets of their conspecifics' actions. That is, the meerkat's being primed to respond to the call by seeking out a bolt hole or digging one does not represent the predator that the call itself is a response to by using the call as a stand-in or proxy for the predator. So, on this account, meerkats' call-based predator-avoidance behavior is not in the first instance a function of their representing the call as standing in or going proxy for something to be avoided or, say, hiding from a predator. Rather, such behavior is in the first instance a function of appropriately representing and exploiting targets or opportunities for action where this usually involves seeking out or creating further targets or opportunities. In the case under consideration, an alarm call is *de agendo* represented as respondable-to by looking for a bolt hole or for a place that, in turn, is *de agendo* represented as diggable. The bolt hole, in either case, is then *de agendo* represented by the meerkat as hideable-in. If the meerkat has a *de agendo* representation of the predator at all, it is one that represents the predator as hideable-from or avoidable.[33] It will have such a representation if it itself comes into perceptual contact with the predator as a target of that *de agendo* representation or through the sentinel meerkat's call acting as an extension of the other meerkat's perception. In the latter case, the sentinel meerkat itself stands in or goes proxy for the

[33] One might worry that this characterization of meerkat behavior is either myopic or threatens an indefinitely open-ended disjunctive set of possible *de agendo* representations for predators. If the characterization is myopic then my presentation of Springle's framework fails to be sufficient as a characterization. If it leaves open indefinitely many possible *de agendo* representations for predators corresponding to each response that a meerkat might have to a predator or sentinel call then my presentation fails in the same way that earlier behaviorist accounts failed. Ethologists will abandon the account in favor of one on which meerkats have internal mental representations that are not directly tied to action. In response, I admit that the example is overly narrow by focusing on only a couple of possible responses to a certain sentinel call. Yet, I maintain that the foraging meerkat's representation of the predator is still as hidable-from or avoidable. Avoidance is the solution to the problem of the predator. And the solution systems organized under that solution can be understood as distinct ways of representing the predator as hidable-from or avoidable. Indeed, this will be the case even if the foraging meerkat merely monitors the predator or takes up the role of sentinel in making a further call to others. In the first case, the meerkat is maintaining vigilance for the predator's action so as to inform further responses. These further responses will include hiding itself, signaling others, going back to foraging if the threat passes, and so on. In the second case, the meerkat is representing the predator as something for conspecifics to hide from or avoid. Of course, this explanation implies that Springle's framework requires other-representation of conspecifics in nonhuman animals. The important point for now is that, however complex the behavior, the core representation is that of the predator as hidable-from or avoidable.

other meerkats but the other meerkats do not represent the call or the sentinel as standing in for anything. The foraging meerkat, then, has *at least* three distinct, yet, strongly linked *de agendo* representations: that of the call as respondable-to by searching for a bolt hole or digging location, that of the bolt hole or digging location as burrowable-into or diggable (respectively), and that of the predator as hideable-from or avoidable. The last of these is the core representation that is most likely to be of interest to ethologists. The foraging meerkat does not have a discursive representation of the sentinel meerkat's call.

In humans, when activated, capacities for unmediated action also constitute embodied modes of presentation of targets. I take there to be three differences between human and nonhuman animal *de agendo* representation. One difference is that, for humans, targets of *de agendo* representations fall under two modes of presentation at once. That is, targets are presented to us in perception both in terms of potentiated actions and objectively as what they are. A second difference is that, for humans, both capacities for action and actions themselves are presented objectively in *de agendo* representations as ways that any appropriately situated agent could at least in principle do what the agent is doing by taking the relevant means.[34] That is, such representations represent "an order which is [*really*] there whenever actions [of the relevant type] are done with intentions" *as such* (Anscombe 1957/2000: §42). Finally, for humans, the needs that a capacity for action functions to satisfy are likewise presented objectively in *de agendo* representations as grounds for practical reasons. For instance, the need for nourishment is represented in the tokening of *de agendo* representations in human beings as providing reasons, say, for preparing a meal, ordering take-out, going to the grocery store, hunting, eating whatever already edible food is readily available, and so on. While nonhuman animals are responsive to such reasons by virtue of the tokening of their *de agendo* representations, nonhuman animal *de agendo* representations do not present those reasons as reasons. Rather, they present needs only in terms of the interactions that function to satisfy those needs.

All that said, *de agendo* representations are not themselves descriptions even in the human case. Rather, they are the priming of those very well-ordered intentional actions that in turn prime further well-ordered intentional acts of deliberating, expressing an intention, or describing, explaining, or assessing those very actions in light of the relevant practical reasons. The latter acts are discursive. They are distinct acts of constructing and using *de substituto* representations for the relevant order. *De agendo* representations represent objectively in

[34] Note that representing one's capacities for action objectively does not require that any other agent at any point in the past, present, or future actually be appropriately so situated or that any other agent at any point actually possess the relevant capacities for action.

human beings in virtue of presenting their targets as what they are, in virtue of presenting their needs as generating practical reasons, and in virtue of themselves being potential realizations of means-end orders that any appropriately situated agent could at least in principle realize. Such orders are the referents of intentional action concepts and those concepts are only given full application in corresponding *de substituto* representations. That is, human *de agendo* representations objectively represent their targets, needs, and the means by which the agent acts on those targets in pursuit of the relevant end(s). Targets are represented objectively by drawing on conceptual capacities involved in perceptual judgment. Needs and actions are represented objectively by drawing on conceptual capacities involved in deliberation as well as the expression of intention and action description, explanation, and assessment. Such "drawing on," in turn, amounts to i) sharing a referent (the relevant means-end order) with the *de substituto* representations constructed and used in deliberation and the like and ii) the transformative effect that tokening those *de substituto* representations has on structuring capacities for realizing that referent. Such structuring transforms human *de agendo* representations from those enjoyed by nonhuman animals into the activation of intentional action concepts *proper*, where such concepts have a form expressible in language as "I/One can do *B* (by doing *A*)." Such transformation puts us in a position to represent our actions and capacities for action as well as that of others not just as targets of other *de agendo* representations but also as targets of *de substituto* representations.

De Substituto ("about that which is substituted" or "about the thing a stand-in or proxy stands in or goes proxy for") representations are activations of capacities for mediated actions.[35] At least paradigmatically, such representations are themselves products of intentional actions that function to solve needs-based problems having to do with limitations to unmediated practical access. For instance, say an orchidologist needs to learn to identify some species that only grows in a remote part of the world that she has not had the chance to visit. A photo of a member of the species can stand in or go proxy for seeing an actual flower. By studying the photo, the orchidologist can learn how to identify the species without ever directly encountering it (a written or spoken description could also do the relevant work, though perhaps less effectively). Of course, the mere existence of a discursive representation will not make it the case that one has mediated practical access to what it represents. Discursive representations can provide mediated practical access only when agents have unmediated practical access to their vehicles. Without at least some unmediated practical access, an agent could not intentionally

[35] Hereafter, I will refer to *de substituto* representations as discursive representations, activities involving the construction and use of *de substituto* representations or the potentiation of such construction and use as discursive activities, and the capacities for such construction and use discursive capacities, discursive solution systems, and discursive skills and habits.

interact with a discursive representational vehicle.[36] According to Springle, it is for this reason that discursive representations have "derived" intentionality (Springle and Buccella 2024: 13). The intentionality of discursive representations is derived from the intentionality of the *de agendo* representations for construction and use of their vehicles constitutive of the construction and use of those discursive representations.

Both *de agendo* representations (or "intentional action acorns") and discursive representations have constitutive appropriateness conditions. The appropriateness conditions of *de agendo* representations are the same as those of the intentional actions of which they are potentiations. That is, the appropriateness of an intentional action is the applicability of that (part of the) solution system for which the agent is accountable relative to the need- and target-facts constituting the system's domain of application on the relevant occasion. Because they are the acorns of indirect intentional actions on targets to which the agent has mediated practical access through direct actions on discursive representational vehicles, the appropriateness conditions of discursive representations envelope their constitutive veridicality conditions (Figure 1.6). In order to gain mediated practical access to the orchids that she is interested in studying, the orchidologist must capitalize on the mapping relations that the photo bears to those orchids and does so through capitalizing on her unmediated practical access to the photo, that is, through using it appropriately as a stand-in or proxy. The social praxis of nature photography is in part dedicated to the production of photos that bear these mapping relations to their objects. It is thanks to this social praxis and the nature photographer's participation in that praxis in taking the photo that the photo bears the relevant mapping relations to the relevant orchids such that the orchidologist enjoys mediated practical access to the latter through her unmediated practical access to the former. That is, the veridicality conditions on the appropriate use of the photo by the orchidologist are partly grounded in the social praxis of nature photography. The same goes for propositional and map-like representations. All three kinds of representation are kinds of discursive representations.[37]

---

[36] A single representational vehicle typically constitutes a multiplicity of discursive representations corresponding to all the distinct actions it potentiates (where the relevant actions may be of different types). On Springle's account, what are often labeled as representations in philosophy of mind are things that function to be potentials for types of action rather than themselves being potentiations of particular actions (Springle refers to these as "concepts," while representations are applications of concepts).

[37] It's important not to confuse *de agendo* representations with either directives or imperatives. According to Springle, both directives and imperatives are discursive representations that represent by standing-in or going proxy for things, namely, actions and situations. They have constitutive veridicality conditions, only understood as an aspect of their appropriateness conditions. The different ways *de agendo* and discursive representations are about the things they represent is reflected in their success conditions. Since *de agendo* representations do not represent by standing-in or going proxy for anything but, rather, by being potentiations of actions that are themselves intentionally directed at targets, they have constitutive appropriateness conditions but not constitutive veridicality conditions. By contrast, directives and imperatives have

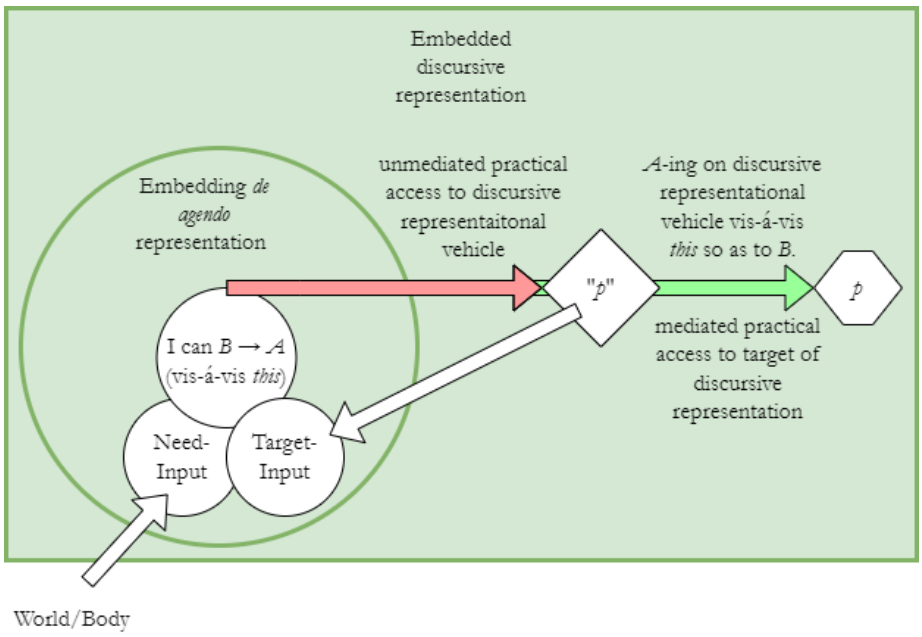

Figure 1.6 Discursive Representations

A diagram depicting a discursive representation. Green coloration represents activation of the embedding *de agendo* representation sufficient to perform an action of discursive representing (in this case, of *p*). I will explain the location "I can *B* → *A* (vis-á-vis *this*) in Section 1.3.2.

As we'll see, deliberation, the expression of intention, and the description, explanation, and assessment of action are all instances of constructing and using discursive representations of means-end orders where such construction and use are themselves completed exercises of corresponding discursive skills that involve ordering the relevant means and ends. Activation of these discursive skills are, in turn, *de agendo* representations for the construction and use of such discursive representations. Expression of intention as well as action description, explanation, and assessment function to solve a variety of problems including instructing another, informing another, getting out of trouble, getting someone off your back, and so on. Deliberation is more specific: it functions to solve the problem of having a need that cannot be satisfied by application of the battery of solution systems that the agent currently has. That is, deliberation is a response to the need to find a way to satisfy some other need. Moreover, once activation of such construction has occurred, one can then turn other discursive

constitutive veridicality conditions. This is because the appropriateness of the indirect actions that they potentiate depends on how they stand-in or go proxy for a target and on whether they do so "faithfully" or "truly" in the ways required for the indirect action to non-accidentally achieve its end. The truth-conditions of an order or request are that the action ordered or requested is to be done. Orders and requests are successful if what they represent as to-be-done is faithful or true to what, in fact, is to be done. Directive and imperative contents, then, are just discursive representations that are parts of acts of (trying to) get someone to make it the case that either the world and the discursive content match—say, if the representation is a picture or image—or make it the case that the world satisfies the content—say, if the representation is a descriptive sentence.

skills onto deliberation, tokening *de agendo* representations for constructing and using discursive representations of one's *de agendo* representations for constructing and using discursive representations of the relevant means-end orders. I take it that an expression of an intention to deliberate, say, "I am figuring out how to do *B*," is an instance of such a meta-exercise of these other discursive skills. And one can iterate the exercise of these latter discursive skills in descriptions or explanations such as, "I was just telling Colin that I am trying to figure out how to do *B*." Likewise, one can iterate deliberation itself in, say, trying to figure out how best to find a solution to some other problem.

Springle's practical access analysis of representation has at least three advantages. First, it is a unifying analysis. The account grounds mental representation understood as *de agendo* representation and public representations understood as discursive representations in the same generic function of potentiating intentional actions. The difference between these kinds of representation consists in whether they are potentiations of unmediated or mediated intentional actions, respectively. Second, propositional attitudes can be accounted for in the analysis in terms of composite *de agendo*-discursive representations and the discursive activities in which they figure. Discursive social praxes employ discursive representations. Humans like our orchidologist have capacities to respond to and construct these representations in appropriate ways and on appropriate occasions. Human capacities to construct and use discursive representations yield *de agendo* representations that are composite: they "embed" the discursive representations they are potentiated actions for constructing and using. For instance, the nature photographer's snapping a picture of an orchid involves his having an embedding *de agendo* representation for producing a photo that accurately captures the flower's features and the photo-to-be-produced is embedded in his potentiated action for snapping that photo. And our orchidologist's potentiated use of the photo is likewise an embedding *de agendo* representation that embeds the photo's discursive content, namely, those objective identifying features of the orchid captured by it and by which she might later come to recognize similar orchids as being of the same species in the wild. It is for these reasons as well that application of intentional action concepts in the formation of truth-conditional representations representing corresponding intentional actions are distinct from those concepts' being drawn on in the activation of the relevant parts of an agent's non-discursive skills or incorporating habits. The former are embedding *de agendo* representations. The latter are non-embedding, yet, rationally transformed *de agendo* representations that merely activate those embedding *de agendo* representations. The former correspond to mental actions, while the latter correspond to non-mental actions (more on this in chapter 2).

A third advantage of Springle's account brings us back to its original purpose, namely, providing a naturalistic, non-reductive, action-first account of perceptual representation primarily in nonhuman animals the intentionality of which is grounded in the developing unified teleological process of intentional action. In particular, Springle's account aims to make sense of two seemingly conflicting facts about perceptual representation. Namely, both humans and nonhuman animals perceive, yet, human perception is unlike nonhuman animal perception in that it provides us objective knowledge of particulars in our perceptible surroundings. Many philosophers assume that the fact that perception is a capacity we share with nonhuman animals means that, since animals seem to lack discursive skills that function to apply concepts, perception in humans must be nonconceptual. But this creates trouble for making sense of perception's epistemic role (McDowell 2008).[38] That is, it is difficult to see how we could ever be in a position to apply concepts like ORCHID or APPLE correctly in ascribing the corresponding objective properties to their proper objects when some such objects are in our presence if our perception of them did not already draw on capacities to apply those concepts. Indeed, it is difficult to see how we could ever rightly judge "that is an apple" on the basis of looking at an apple if our perception is, like that of animals, thoroughly nonconceptual. Maintaining that human perception is nonconceptual creates a gap between perception and conceptual content that draws on perception. Many philosophers assume that this gap can be easily avoided. After all, perception enters into explanations of broadly intelligent behaviors in animals where it plays a broadly epistemic role in providing a kind of action-guiding awareness of the world. The problem, however, is that unless perception's epistemic role in nonhuman animals involves conceptual capacities after all, the gap has not been addressed. Animal's perceiving targets in terms of interactions they can generate will not, on its own, give us the content required for perceptual judgment. Those who take overcoming the gap seriously propose a different solution: perception is a psychological genus of which human and nonhuman animal perception are distinct species (Marcus 2012, 2021; Boyle 2016, 2022).

The third advantage of Springle's account is that it provides a way of articulating this genus-species proposal that captures the behavioral expressions of perceptual awareness in nonhuman animals without attributing superfluous discursive skills to them and that allows that perceptual states in humans are conceptual in virtue of having composite elements. Animal perceptual contents are non-conceptual activation of capacities

[38] McDowell (1994), echoing Sellars (1956), labels the trouble involved in claiming that perception either brings conceptual content with it immediately in the complete absence of discursive activity or that it is thoroughly non-conceptual despite being able to perform its traditional epistemic role "The Myth of the Given."

for interaction with perceptible objects through and through. That is, animal perceiving is thoroughly *de agendo*. Their broadly intelligent behaviors are a function of the appropriateness of the capacities for action that are activated in the relevant circumstances. But they do not have nor can acquire discursive skills and, so, are not in the business of categorizing or otherwise cognizing the objective features of particulars in their perceptible environment. By contrast, human perceptual contents are embedding *de agendo* representations that embed discursive representations. The discursive components render human perception conceptual so that such perception can play discursive roles, say, in justifying perceptual judgments. Perception in humans, then, just is the activation of a range of solution systems for the construction and use of discursive representations. We can avoid creating a gap between perception and conceptual content that draws on perception while respecting the fact that both humans and animals perceive. Perception, as a genus, is a capacity for unmediated practical access. Perceiving is always a matter of a creature's having *de agendo* representations. But in humans, it is not *only* a matter of perceiving targets in terms of interactions that we can generate. For perception in humans is rationally transformed: it incorporates the discursive form of practical access. And such transformation is perfectly natural by virtue of being a kind of second nature.[39] That is, we come to have embedding *de agendo* contents in perception thanks to our inculcation into a variety of social praxes for making rational use of such contents—where that use entails latching onto the objective features of things in our perceptible environment that those contents represent—and thanks to our having needs which call on the application of solution systems for constructing and using corresponding discursive contents.

The rational transformation of capacities for action into rational dispositions is a bit different. For, as Small (2012: 218-221) points out, the distinction, as articulated by McDowell (1994), between active application of concepts in judgment (perceptual or otherwise) and the passive drawing on of concepts in transformed perception is a non-starter. There is no *passive* drawing on of concepts in the performance of rational action. So, Small advocates for a different distinction. Intentional action concepts are "actively synthesized" in deliberation or are "immediately instantiated" in the successful performance of intentional

[39] Although both Springle's account of mental representation and my use of it in explicating intentional action concepts are pitched entirely at the whole-organism level and reserve an essential role for social praxes, neither is hostile to empirical inquiry into the mind. For among such inquiries are those conducted by the so-called "social sciences," among whose objects are the very praxes to which both Springle and I appeal. What must be stressed is that empirical inquiry into the neurophysiological underpinnings of mental representation and intelligent agency as they exist in individuals will not give us complete accounts of their intentionality or their nature.

actions (Small 2012: 222-223). Immediate instantiation is a matter of making true the means-end order that one is in a position to actively synthesize in deliberation without having to actually deliberate.

With the addition of Springle's account, we can identify immediate instantiation with activation of the relevant part(s) of some skill or incorporating habit sufficient to initiate performance rather than with actual realization of the relevant means-end order in ongoing and successful performance. This is important because any account of intentionality needs to be able to explain misrepresentation. And if we were to identify the drawing on of intentional action concepts with realization of a means-end order in performance then we could not explain misrepresentation. For, in that case, we would be stuck saying that an agent could only represent herself as doing *B* by doing *A* when she succeeded. Identifying immediate instantiation with activation sufficient to initiate performance rather than full manifestation in ongoing or successful performance gives us a way to account for misrepresentation in practical cognizing. Namely, such misrepresentation occurs when the intentional action concepts that are drawn on in activation of the relevant part(s) of the relevant skill or incorporating habit are not appropriate for satisfying the relevant need. That is, the means-end order is represented as satisfying a need in virtue of such activation, but, were she to have that need and perform the action, that need would go unsatisfied.

Appealing to Springle here also gives us a way of accounting for intentions without identifying them with truth-conditional propositional attitudes modeled after belief. That is, on the powers framework as I am elaborating it, intentions are just the activation of the relevant part(s) of some skill or incorporating habit. Construction and use of a discursive representation of the relevant order as achieved by deliberation is a way of forming intentions because such construction and use function to activate the relevant part(s) of those dispositions. All that said, I take this emendation to be in keeping with Small's contributions to the powers framework. After all, he likewise treats intentional actions as teleologically unified developing processes whose *a priori* moments constitute stages of those processes. Springle's insight is to identify the potentiation of intentional actions as among those moments. Doing so gives us a better foothold for the intentionality of intentional action concepts.

Returning to the rational transformation of capacities for action, such transformation is nonetheless connected to the rational transformation of other mental faculties in human beings in that the former allows formation of embedding *de agendo* representations for the construction and use of corresponding embedded discursive representations specifically of the means-end order to be instantiated. It is important to note that such formation is merely the activation of the corresponding discursive solution systems and not their

application. For holding that application of a solution system for "doing *B* by means of doing *A*" is at the same time an application of the solution system for discursively representing that order runs the risk of running together intentional action and deliberating about that action, expressing the corresponding intention, or describing, explaining, or assessing that action. Rather, a product of the rational transmission of skill is that potentiation of intentional actions in human beings itself potentiates intentional actions of deliberating, expressing the corresponding intention, or describing, explaining, or assessing the former actions. Rational transformation of capacities for action into skills, then, is a matter of structuring those capacities themselves such that they are internally calculatively structured and such that their activation and application merely activates corresponding discursive skills (Figures 1.7 and 1.8). Again, the means-end orders embodied in the agent's skills are objective in that they represent "doing *A*" as a way of "doing *B*" for anyone with the relevant solution system relative to its domain(s) of application. Such orders, so embodied, are the agent's intentional action concepts. That is, these embodied means-end orders are objective representations of the form "I/One can do *B* (by doing *A*)." And the agent's being in possession of such embodied means-end orders, then, is her possessing the relevant intentional action concepts. When activated in the initiation of an agent's intentional action, such embodied means-end orders are objective representations of hers of the form "I am doing *B* (by doing *A*)." They pick out "an order which is [*really*] there whenever actions [of the relevant type] are done with intentions" *as such*, regardless of the stage in the life of that action (Anscombe 1957/2000: §42). Picking out that order in action is an agent's immediately instantiating and, so, drawing on those intentional action concepts through mere activation of discursive skills for deliberation, expression of intention, and action description, explanation, and assessment.

All that said, the objective means-end order captured in an agent's intentional action concepts is the same order represented discursively in her deliberation, in the expression of her intention, in the description of her own or another's corresponding action, in the explanation of that action, and in the assessment of that action. The content is not shared by virtue of having the same vehicles, that is, discursive representational vehicles. Rather, it is shared between the expression of intention, description and the like of one's own action, and activation of the relevant part(s) of the network(s) constitutive of one's skills by virtue of having the same referent and by virtue of having the same mode of presentation. That is, the action so ordered is known by the agent in all three of these cases as what *she* is doing, as an order that she is realizing in action. This mode of presentation is an essential feature of the activation of the relevant part(s) of her skills or incorporating habits.

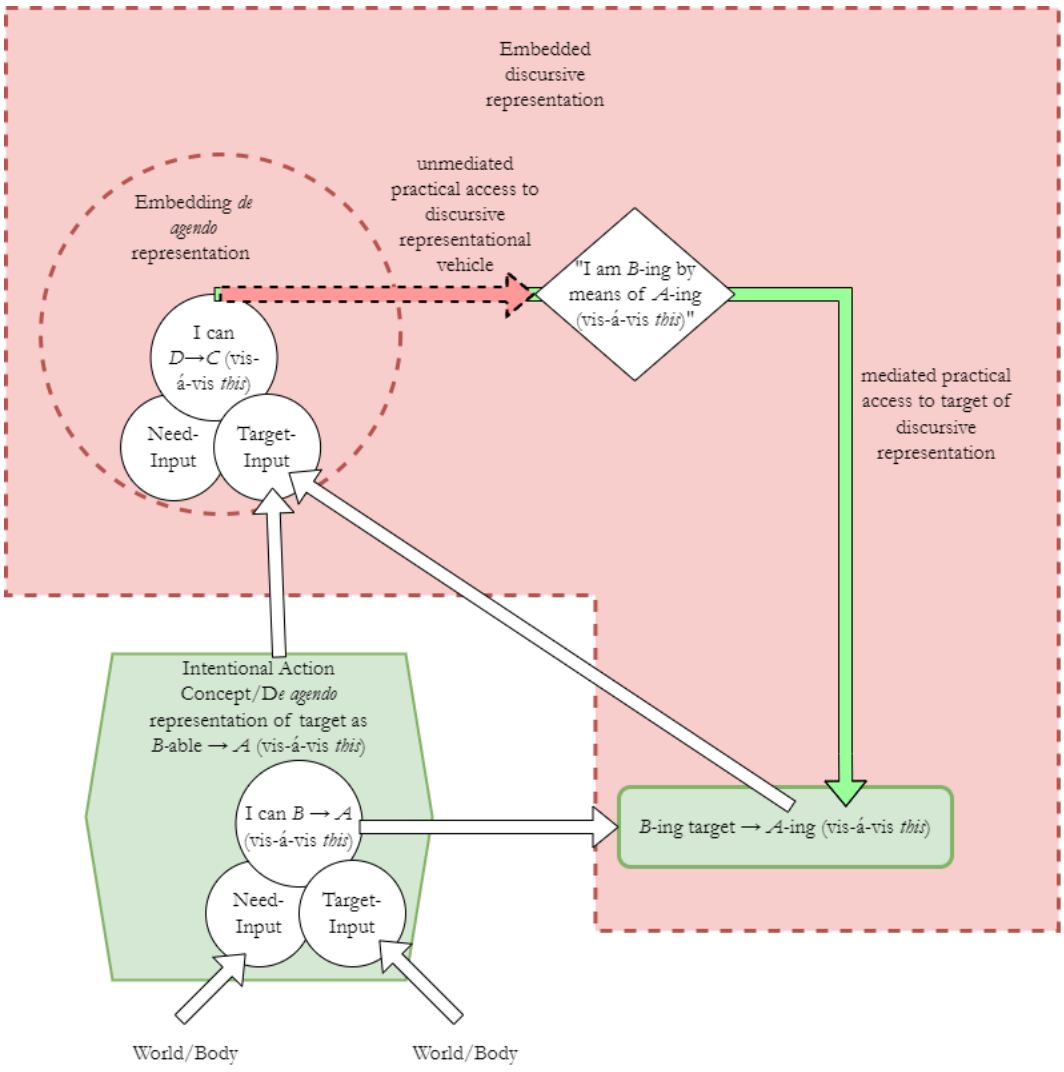

Figure 1.7 An Intentional Action

A diagram depicting an intentional action as performed by a rational agent. Red coloration represents activation in the absence of performance. Green coloration represents activation sufficient for performance. In this case, an intentional action concept is fully activated and, thus, develops into the performance of an intentional action of doing *B* by doing *A* (vis-á-vis *this*) (more on this locution in Section 1.3.2). The diagram depicts a capacity for action that has already been rationally transformed; that is, it depicts the manifestation of a skill. The diagram can be misleading in that I have not depicted the complex structure of skills. Skills consist in intentional action concepts and discursive sub-skills that form an interconnected network related to each other as means and ends. This is not depicted here. To get a sense of what the complex structure of a skill looks like using the diagrammatic conventions I have used here, I ask that the reader imagine a network of interconnected hexagons and gnomons (see Figure 1.8). The top of the diagram and the arrows indicating drive of target-input from the activated intentional action concept and the action to the embedding *de agendo* representation illustrate the powers framework's commitment to the claim that both such activation and the manifestation of a corresponding skill prime the agent to deliberate, express her intention, or describe, explain, or assess what she is doing.

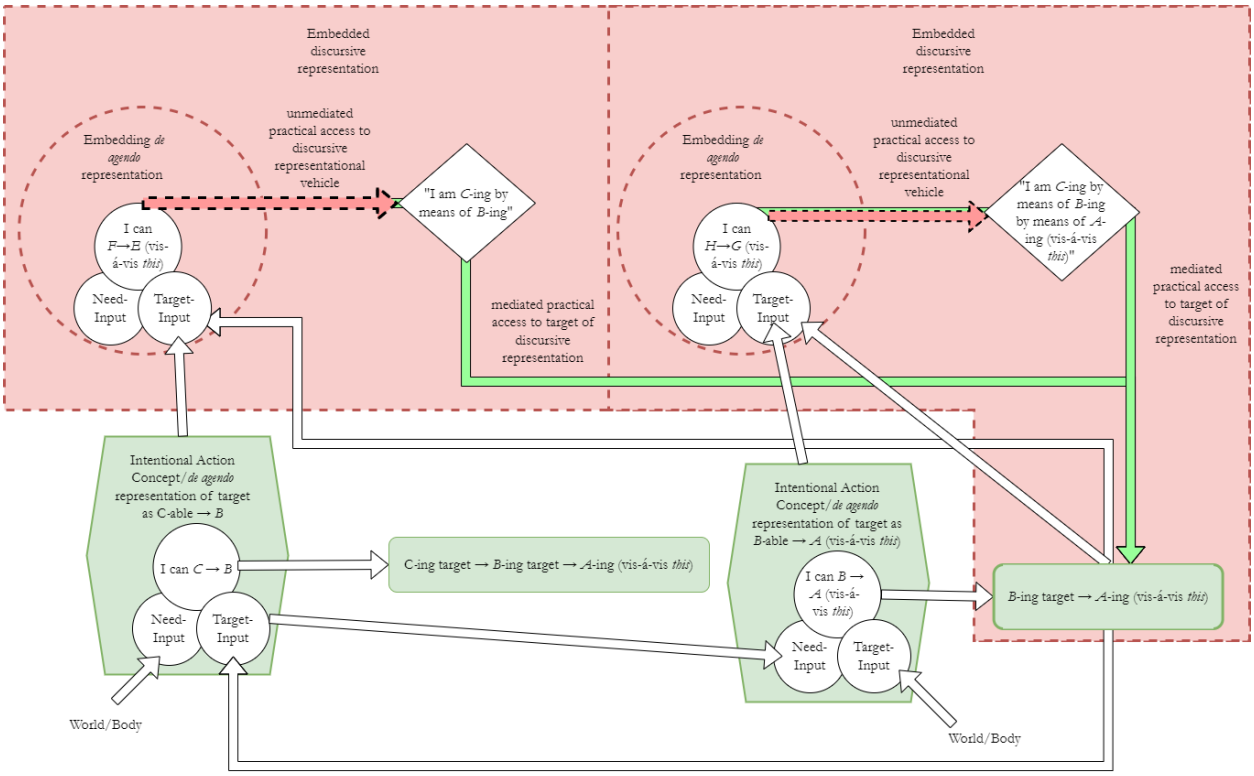

Figure 1.8 An Intentional Action Involving Multiple Steps

A diagram depicting the performance of an intentional act of doing *C* by means of doing *B* by means of doing *A* (vis-á-vis *this*) in the absence of deliberation, expression of intention, or action description, explanation, or assessment. Red coloration represents activation in the absence of performance, while green coloration represents activation sufficient for performance. What is depicted is the human analogue of the non-human animal's use of a sub-solution to discover or generate a target for the application of the solution that subordinates the sub-solution. The reason I am depicting it to give readers a sense of the complex structure of skills. Skills are dispositions composed of calculatively connected intentional action concepts. So, a complete diagram of a skill would consist in many hexagons and gnomons connected to each other in much the way depicted here. The only missing element would be modulatory intentional action concepts.

The content representing the objective means-end order is shared between the expression of intention, description and the like of one's own action, activation of the relevant parts of one's skills or incorporating habits, and deliberation in virtue of having the same referent and by virtue of all four involving or simply being the activation of those parts. Again, activation of those parts primes application of discursive sub-skills whose function is to construct and use discursive representations that apply those same intentional action concepts. It is that activation in both the initiation of intentional action and in deliberation that makes it the case that the former puts the agent in a position to do the latter. Finally, the content is shared between these four and the description and the like of *another's* action for the same reason that it is shared between the prior three and deliberation. The differences here are the mode of presentation and the kind of reason or knowledge involved. Whereas deliberation is practical and presents the order objectively as well as one that the agent herself can

realize, the description and the like of other's actions is an expression of speculative knowledge and presents the order objectively as well as one that someone else is realizing.

How an agent comes to have skills and incorporating habits with internally calculative structure that reflects in principle publicly available ways of doing things is in the first instance a function of the role that corresponding discursive representations have for internally organizing those dispositions as part of the rational transmission of skills. Such organizing is a key part of rational transformation. I am not going to give a full account of such transformation here. I am assuming it and elaborating it to the extent that it aids in distinguishing the powers framework from its causalist rivals. However, part of the story involves a deeper understanding of practical reason understood as something that includes but is not fully exhausted by deliberation. Practical reason so understood backs cogent practical cognizing and, with it, practical knowledge. To start, I will discuss what I call "practical want" before moving on to the form of practical inference and their mutual instantiation in cogent practical cognition.

1.3.2 Practical Want as the Generans of Practical Reason

Thus far, we have an account of the intentionality of intentional action concepts: embodied means-end orders that are structured such that anyone with the relevant solution systems relative to their domain(s) of application also counts as possessing those concepts. These concepts are objective representations of the form "I/One can do *B* (by doing *A*)." And an agent counts as possessing such concepts in possessing such embodied means-end orders as skills or incorporating habits. She counts as drawing on those concepts and, so, representing "I am doing *B* (by doing *A*)" when such dispositions are activated such that she is initiating that action. At this stage, I want to hook the account of the intentionality of intentional action concepts up with practical reason. To start to do this, I draw from §§34-40 of Anscombe's *Intention* (1957/2000) so as to lay out what I call "practical want." I identify practical want with what Anscombe identifies as "wanting" in §36 as possessed specifically by human beings. That said, I will relate each part of the notion to Springle's framework as I go. And, speaking metaphorically, I take practical want to be the "*generans*" of practical reason.

Here's my account of practical want:

Practical Want

An agent practically wants[40] something just in case:
a) she takes that thing to be
    a1) in some way good and

[40] The double entendre here is not lost on me.

a2) in some sense at a distance; and

b) she is trying to get it, that is, she

b1) is engaging in movement or action that (the agent takes) is of use towards getting it and

*b*2) takes that it is gettable by that movement or action.[41]

Other things being equal and assuming she does not change her mind, any time an agent practically wants something and practically reasons from it to means that she can take immediately, she engages in intentional action. Let's unpack practical want. First, the sense of "good" as used in (*a1*) is "not that of what is really [or ethically] good but of what the agent conceives to be good; what the agent wants would have to be characterizable as good by him, if we may suppose him not to be impeded by inarticulateness" (§40). That is, "*[b]onum est multiplex*: good is multiform" and a requirement on practical want is that an agent capable of communicating must be in a position to characterize what she wants as falling under some good (§39). Following Volger (2002) and Schwenkler (2019: 142-143), I take it that Anscombe's multiform good admits of three general categories: the useful, the befitting, and the pleasant (*cf.* 1957/2000: §35).[42] Importantly, what Anscombe is pointing out here is that the agent can be wrong about whether what she is aiming at falls under the good that she takes it to fall under and, yet, her characterization of it as desirable is nonetheless intelligible. This should be unsurprising: we find a belief that happens to be false intelligible when expressed by someone holding it. This is simply the practical analogue. Second, in reference to (*a2*), something can be at a spatial, temporal, or conceptual distance from the agent (§41). Sending someone a gift through the mail is a way of closing spatial distance between oneself and one's end of giving that person a gift. Having coffee in preparation for an all-nighter is a way of closing temporal distance between oneself and one's end of staying awake all night. And coming to a decision is a way of closing conceptual distance between oneself and one's end of figuring out what to do. "Distance," then, is meant only to signal that, in order to intelligibly practically want, the agent

[41] I formulate (*b2*) this way due to what Anscombe says immediately before about the features of animalistic wanting as exhibited in the primitive sign of trying to get, namely, "movement towards something, and knowledge, or at least opinion, that the thing is there" (§36). Because her introducing "the idea of the thing" is supposed to cover cases of wanting non-existent things by humans, there is no point in adding "as being there" to it. The thing does not exist and the agent knows this. So, her idea of it will not be one of it being anywhere. Yet, without this context, it can seem as though all that is required for the second feature of wanting in the human case is the idea of the thing *full-stop*. And thinking this would be a mistake, since then one could intelligibly try to get things that one knows they cannot get. With context, it is clear that "the idea of the thing" is meant to include the thing's being gettable by the agent.

[42] Part of what this means is that, according to Anscombe (and Aristotle and Aquinas), someone professing to want something and who says that he wants it because it is good but who genuinely does not consider it useful, befitting of them in any way, or pleasant is thus unintelligible.

must not think that she has already achieved her end. In Springle's framework, (*a*)'s obtaining is a matter of the agent's registering the obtaining of both a need-fact (*a1*) and target-fact (*a2*) such that she is primed to apply the relevant solution system (an intentional action acorn is produced). (*a1*)'s obtaining then amounts to activation of the Need-Input sub-capacity, while (*a2*)'s obtaining amounts to activation of the Target-Input sub-capacity. If the agent is capable of communicating, her being so primed in turn provides a target-fact within the domain of application of her discursive solution systems for deliberating, expressing her intention, or describing, explaining, or assessing her action. That is, the action itself becomes a potential target for deliberation or communicating her intention in such an agent upon that action's potentiation.

The movement or action referred to in (*b1*) can include mental movements as well as non-mental ones. That said, such movement or action is not thereby intentional action despite in some instances being supposed by the agent to be of use towards getting what she wants. First of all, some such movements or action can close (some of) the distance between an agent and what she wants without her being aware of it. That is, in the absence of satisfying (*b2*), an agent can satisfy (*b1*) with movements or actions that she does not realize close the distance between herself and what she wants. As Oedipus starts uncovering his own past, his investigating satisfies (*b1*) with respect to finding the murderer of Laius. Suffice to say, he does not take himself to be doing this. Rather, he takes himself to be doing things that close the conceptual distance between his current epistemic state and one in which he knows his past. Second, even if the movements or action are taken by the agent to be of use towards getting what she wants and she knows this, the agent might be plain wrong. In which case, if she satisfies (*b1*), she does so without doing whatever it is she takes herself to be doing (*cf.* §23). She is just lucky in actually closing the distance.[43] *A fortiori*, she is not doing what she takes herself to be doing intentionally.

In the absence of satisfying (*b1*), the agent's take on her movement or action referred to in (*b2*) is consistent with one's doing nothing to get what one wants. Indeed, an agent can satisfy (*a1-2*) and (*b2*) and, so, want something without thereby being compelled to act. In such cases, she is motivated to act but does not act on that motivation. These cases include those in which one experiences "the prick of desire at the thought or sight of an object [where nonetheless one] then does nothing towards getting the object" (§36). In Springle's framework, such a thing occurs when some need is driving potentiation of some action more than the need here registered by the agent in having the "prick of desire." Application of the solution system activated by that

[43] Here is a case from Small (2012) of what I have in mind: "Suppose, for example, that to operate [a] pump, which unbeknownst to our man does not work, he has to stand on a platform, and that in doing so he has activated the automatic pump, of which he is not aware, and is thereby replenishing the water supply [and thus doing what he intends] for as long as he stands on the platform" (161).

desire is thus suppressed. In the good case, such suppression most likely occurs because doing something else better maintains the superstructure of solution systems. The agent herself might event realize this, say, in saying or thinking "better to save my money for when I need it than spend it on this dress even though it would suit me well" (*cf.* §33). Indeed, rather than such wanting in the absence of action manifesting some inconsistency in one's practical rationality, the suppression involved here and possibly recognized by the agent is a manifestation of consistency in her practical rationality. Moreover, it is a kind of consistency in practical rationality that nonhuman animals can exhibit.

Finally, (*b1*) includes the parenthetical that the movements or action in question are taken to be of use towards getting what is wanted and (*b2*) is disjunctive in order to include cases where the agent tries to get something that is in fact not gettable either by the relevant movements or actions or at all. For instance, consider that trying to get a key that lies on the other side of a locked door can consist in pulling at the door handle, wedging something in the door to try to disengage the lock, and so on to no avail (*cf.* §36). The agent cannot get what she wants, namely, to obtain the key. Nonetheless, she still counts not just as wanting it but as *practically wanting it*. And this is because she takes her pulling, wedging, and so on to be of use towards getting the key and because she takes it that the key is gettable by these means. In Springle's framework, (*b*)'s obtaining is a matter of (*b2*) the activation of a solution system by some target-fact within its domain of application and (*b1*) the agent's applying that solution system successfully or unsuccessfully.[44] Again, if the agent is capable of communicating, the performance initiated in her application of that solution system in turn provides a target-fact within the domain of application of her discursive solution systems for deliberating, expressing her intention, or describing, explaining, or assessing her action.

There are two final clarifications I want to make. First, I take it that the agent's take on things referred to in (*b2*) is not such as to require positing a corresponding truth-conditional propositional attitude. Second, I

[44] On this way of framing practical want, an agent can only satisfy (*b1*) as part of practically wanting *something*. She cannot satisfy (*b1*) in the complete absence of satisfying (*a1-2*) and (*b2*) with respect to everything that she takes as falling under some good, being at some distance, and being gettable by the relevant movements. She can satisfy (*b1*) in the absence of satisfying (*a1-2*) and (*b2*) with respect to some target. In such a case, the agent is doing something that she does not take to be good or at a distance, that she does not take it will close that distance, or some combination thereof. The Oedipus case is one of satisfying (*a1-2*) and (*b1*) in the absence of satisfying (*b2*). Cases of foreseen consequence and *akrasia* would both be ones in which the agent satisfies (*a2*) and (*b1-2*) in the absence of satisfying (*a1*). Cases where the agent keeps trying to do things that achieve a goal she has already *knowingly* achieved are ones in which she satisfies (*a*1) and (*b1-2*) in the absence of satisfying (*a2*).

take it that, likewise, the agent's conception of the thing wanted as laid out in (*a1-2*) is not such as to require positing a corresponding truth-conditional propositional attitude.

Starting with (*b2*), I take it that the agent's representing what she wants as gettable by her movements or action does not require that she be consciously entertaining a corresponding proposition at the time of acting in order for her to count as trying to get it. There are at least three reasons for this. First, committing to there being some such state begs the question against the powers framework and in favor of its causalist rivals regarding their ability to explain intelligent agency in general, let alone mental action. Secondly, given the adoption of Springle's account of intentionality, the powers framework holds that such a conception can be instantiated in the potentiation and performance of intentional action itself without the need to posit any additional mental states alongside performance. Consider the master pianist. Her intentional act of playing *Wrong Note* on some occasion is an exercise of her practical knowledge, say, that her moving her hands thus-and-so, in these circumstances, just is her playing the opening passage *leggiero* and *scherzando*. If this is right then there is no need to posit an additional truth-conditional propositional attitude coincident with performance. Third, requiring such a state unnecessarily and severely limits the explanatory power of practical want. For we should consider as part of our *explanandum* intentional actions that look far into the future. Consider such expressions as "I am trying to get my PhD," "I am going to get my PhD," and "I am getting my PhD." Considered as genuine expressions of intention and of intentional action on the part of an unwavering graduate student over the course of her time at graduate school, they imply that the student practically wants to get her PhD and that it is her practically wanting to do so that motivates and sustains her taking the relevant actionable means to this end. If this is right, then it cannot be a requirement on practical want that the agent's taking it, say, that a PhD is gettable by doing lots of research, writing, publishing, attending of conferences, doing service work, etc. is necessarily realized in a truth-conditional propositional attitude of hers. Rather, her take on the situation can be understood in terms of the potentiation and actualization of the intentional actions which she takes to be efficacious means to getting her PhD and in terms of the potentiation and actualization of discursive acts of deliberating, calculating the former intentional actions as means to her end. No mental representations with truth-conditional contents need enter the picture.

Moving onto (*a1-2*), I take it that the agent's representing what she wants as useful, befitting, or pleasant and as being at a spatial, temporal, or conceptual distance does not require that she be consciously experiencing the prick of desire (or any other conative attitude considered as an occurrent mental state) at the time of acting in order for her to count as practically wanting it. There are at least three reasons for this. First,

committing to there being some such state similarly begs the question against the powers framework. Second, requiring such a state similarly unnecessarily and severely limits the explanatory power of practical want. If she is unwavering, the graduate student embarking on a PhD and completing it will count as practically wanting to get a PhD for the whole time she is in graduate school. Echoing Anscombe echoing Wittgenstein, the conception under which an end is pursued cannot necessarily consist in a single persisting occurrent mental state, for no such state could have the consequences of practical want (*cf*. §40). Third, given that practical want by definition involves movement that is (taken to be) of use towards getting the thing she wants, it is a mistake to require that such a state is supposed to "get the action going," as it were. To again speak metaphorically, the mistake here is to identify what is sometimes *fuel* for the *generans* of practical reason, namely, the agent's subjective experience of registering a need, with the *generans* itself.

This leads us back to how practical want is the *generans* of practical reason. For Anscombe, practical reason as we possess it is our being able to immediately instantiate or actively synthesize "an order which is there whenever actions are done with intentions" (§42). That is, practical reason consists not just in the capacity for deliberation but includes as well our possession of well-ordered intentional action concepts. Following Aristotle, Anscombe claims that practical reason starts from "*τὸ ὀρεκτόν* (the thing wanted)" (§34) where what is wanted is characterized or understood by the agent "as desirable" (§37) and concludes in "an action whose point is shewn by the premises, which are now, so to speak, on active service" (§33). The premises connecting the agent's conception of the thing wanted to the ensuing action are "calculations of means to ends, or of ways of doing what one wants to do" (§38). Thus, practical reason(ing) has the "form of a calculation what to do," that is, it has the form of means-end calculation (§33). Recall that an agent counts as practically wanting something just in case:

a) she conceives of that thing
    a1) as in some way good and
    a2) as in some sense at a distance; and
b) she is trying to get it, that is,
    b1) is engaging in movement or action that (the agent takes) is of use towards getting it and
    *b*2) takes it that it is gettable by that movement or action.

Through (*a1*), practical want furnishes practical reason with an object, namely, the thing the agent wants as taken by her to be useful, befitting, or pleasant. In this way, practical want orients practical reason. In Springle's framework, the obtaining of both a need-fact and a target-fact activates the relevant solution system towards

satisfying the relevant need. Through (*a*1) and (*a2*), practical want motivates and makes intelligible identifying means to close the spatial, temporal, or conceptual distance between her and what she wants. In this way, practical want makes intelligible engaging in practical reason(ing) towards finding and implementing "doing *A* (in order to do *B*)." In Springle's framework, what is activated here is a solution system for generating further opportunities for the immediate application of a further (sub-)solution that the agent possesses. Through (*b1*) and (*b2*), practical want motivates implementing or identifying the *appropriate* means to close that distance. In this way, practical want gives practical reason materials to work with and to shape towards achieving its (minimal) standard of correctness, namely, identifying and implementing means that are actually sufficient to achieving the relevant ends. In Springle's framework, activation of the relevant solution system is itself the implementation of those means or is a searching out either of those means or of circumstances in which a means one is already in possession of can be implemented. Finally, through (*b2*), practical want gives practical reason its particular form of validity, namely, goodness preservation from the end sought to the means the agent immediately instantiates. In this way, practical want establishes that (minimal) standard of correctness for practical reason. In Springle's framework, the sub-solutions the agent ends up applying fall under whatever solution system is appropriate to satisfying the relevant need in the current circumstances. Practical want, then, animates the capacity for means-end calculation towards identifying actually sufficient means to the agent's ends in both intentional action and deliberation. It is in this sense that it is the *generans* of practical reason.

1.3.3 The Form of Practical Reason

I have claimed that intentional action concepts are embodied means-end orders that can be generalized to any appropriately situated agent in possession of the relevant skills or incorporating habits. And I have argued that practical want of something drives practical reason(ing) towards getting it. Now all that is left is to give an account of the form of practical reason and its cogency in both effective deliberation and successful intentional action. To start, we can contrast Anscombe's characterization of practical reason with her characterization of theoretical reason. Practical reason is very different in its content, form, and function (read: *telos*) from theoretical reason (Small 2022). Anscombe (1957/2000) characterizes theoretical reasoning in contrast to practical reasoning in §33 of *Intention* as "reasoning towards the truth of a proposition, which is supposedly shewn to be true by the premises," "an argument that something is true," "reasoning for the truth of a conclusion," and "reasoning from premises to a conclusion which they prove." The starting point of theoretical reason, then, is something already known or believed to be true and such reasoning concludes in a judgment

that something else is the case in light of the thing one already knows or (justifiedly) believes. The premises connecting what is known to the judgment are considerations that prove, show the truth of, or provide evidence in favor of the truth of the proposition that one ultimately judges to be true. The contrasts between the content, form, and function of practical and theoretical reason, then, are as follows. **Object**: things as practically wanted and achievable by the agent versus anything that could be true or false and which she could consider. **Form**: calculation to an immediate action that has its point in getting, doing, or securing what the agent practically wants versus considerations of evidence for or against some conclusion. **Function**: producing action that closes the distance between the agent and what she practically wants versus (merely) producing knowledge of something as being the case. As Anscombe puts it "the mark of practical reasoning is that [**Object**] the thing [practically] wanted is *at a distance* from the immediate action, and the immediate action is [**Form**] calculated as the way of [**Function**] getting or doing or securing the thing [practically] wanted" (§41; original emphasis).

Again, the form of practical reason is means-end calculation, and its function is to calculate actionable means from some end. However, we should not identify practical reasoning as Anscombe understands it with the mental action of deliberating (*cf.* Small 2022: 278-279). For one, Anscombe claims that a proper practical syllogism never contains a premise containing the content "I want…". She clarifies Aristotle's account of practical reasoning by claiming that "the description *under which* [the thing practically wanted] is aimed at is that under which it is *called* [the thing practically wanted]" (§35; original emphasis) and that the thing practically wanted is "characterised […] as desirable" in this description rather than properly described (by the agent) as something wanted (§37). Second, Anscombe clarifies that Aristotle is not describing any "actual mental processes" by which agents reason from something they practically want to the means to getting, doing, or securing it but, rather, that such reasoning "describes an order which is there whenever actions are done with intentions" (§42). It would at the very least be strange to suggest that the mental process of deliberating never involves instances in which the agent explicitly conceives of the thing she is after under "I want…". But, most importantly, deliberating is an actual mental process of calculating what to do. Anscombe's clarifications of Aristotle's position, then, suggest that we should not identify practical reasoning with deliberation. Rather, we should understand it to include not just the capacity to deliberate but intentional action concepts and the means-end orders they constitute as embodied in rational agents' skills and incorporating habits as well.

Small finds materials in Anscombe's later "Practical Inference" (2005) for the claim that the practical reasoning that backs practical knowledge that, in turn, is constitutive of intentional action in general has a distinctive teleology, i.e., end, purpose, "*use* or *function*" (2022: 254; original emphasis). Indeed, he writes:

> The Teleological [sic] conception of practicality begins with the idea that practical reasoning is reasoning from an objective to the implementation of means in intentional action [...] practical reasoning *is for the sake of* realizing one's objectives through action. It is not just that the starting point of a practical inference is called an objective and the conclusion an intentional action: 'like its typical result,' acting with an intention, practical reasoning itself 'is done *with a view to* an end and *on account of my insight* that it helps to bring this end about' (267; original emphasis).[45]

Moreover, in spelling out this teleological conception, Small claims that its effect on the form and object of practical reason is substantial: "practical reasoning has a distinctive subject matter and a distinctive form, but these are both explained by its distinctive use" (254). Finally, he says that, most importantly, the "premises of practical reasoning turn out to be themselves distinctively practical thoughts; comprehending this reveals the internality of perception, skill, and the first person to practical reasoning" (254). In this subsection, I further unpack this last claim with a view to laying out the form of practical reason.

How is it that we close the distance between ourselves and what we practically want through the performance of actions which we calculate will close that distance? Small (2022) takes answering this question as his primary target in a long meditation on the similarities and differences between theoretical and practical reason. He primarily draws from Anscombe's (2005) "Practical Inference" but also draws heavily from Ford's (2016) "On What Is in Front of Your Nose." Since I take Small to be the person to have most recently and most thoroughly developed the powers framework, I will focus on his (2022) and consider it in light of his other work on the framework. Moreover, I later draw insights from the chapter about the relationships between skills, habits, and one particularly important (again) domain-relative skill for rational agents, namely, the capacity for deliberation.

Small (2022), following Anscombe, holds that what distinguishes practical reason as practical is the "service to which [its premises] get put" (2005: 125). This difference in end, purpose, use, function, or *telos* is what makes for the differences in form and object that Anscombe, following both Aristotle and Aquinas, points out as distinctive of practical reasoning. Small (2022) further elaborates the differences this teleological practicality makes to the object and form of practical reasoning in (I count) at least twelve ways (for a full account of them and their relationships to Springle's framework, see Appendix B). I will summarize them as follows: (1) the objects of practical reason are "future contingents affectable by action" (Small 2022: 267). (2) Practical reason relates these future contingents from some end to some means that the agent can immediately perform through what Small calls the "means-connective" and which has the form "*I* can do *B by means of*

[45] Small is quoting Müller (1979: 97) here.

doing *A*." Small claims that the means-connective is "a *formal* concept of *practical* thinking: a practical conditional 'connective' that is a specific determination of [the material conditional]" (268; original emphasis). He introduces the formal notation "→" for the means-connective. So, we can represent the agent's connecting intentional action concepts from some end, say, "doing *B*" to some means she can immediately implement, say, "doing *A*" in practical reasoning as "doing *A* → doing *B*" (274; original emphasis).

(3) Following Anscombe, Small identifies "practicable" means—i.e., those that the agent can immediately perform—as actions that are sufficiently specific relative to the agent and fully particular with respect to their proximal targets (2022: 269). He initially represents practicable means as "doing *A* vis-à-vis *this,*" where *A* is something the agent is sufficiently skilled at doing such that she can do *A* here and now and where "*this*" is whatever fully particularized target(s) the agent is to act on in doing *A* here and now (2022: 268-272). (4) Following Anscombe and Ford (2016), Small holds that practical reason moves from ends that tend to be general and generic towards means with fully particular targets and that are sufficiently specific. For instance, an agent practically reasons from practically wanting to make *some* hut inhabitable to finding means of making *this* hut inhabitable, means which she can implement here and now. (5) Following Ford (2016), Small (2022) holds that perception plays an essential role in practical reasoning in aiding the agent's identification of a practicable means. According to Small, to reach a practicable means that she can implement here and now, the agent must make perceptual contact with particulars whose mode of presentation is that of an opportunity to implement the relevant means, i.e., as a "*this* that I can do" or, what amounts to the same "doing *A* vis-à-vis *this.*"[46] (6) Likewise following Ford, Small holds that the "premise" of practical reasoning identifying practicable means via perception also involves recognizing "obstacles and opportunities" to the performance of the relevant action (2022: 272). Small adds that what counts as an obstacle or opportunity is relative to the agent's skill. As he puts it "ten-knot winds present an opportunity to an advanced sailor but an obstacle to a novice" (2022: 272).

(7) Adding to Ford's (2016), Small argues that an agent's knowledge of any means-connection between intentional action concepts as practicable requires in addition that she "actually believe [that *she* can do *B* by

[46] Because I am trying to make room for mental action on the powers framework and because I wish to allow that some mental actions are fully covert and not necessarily initiated by perception of a cue, I cannot follow Small to the letter here. Yet, Springle's *de agendo* representations can be of help. That is, instead of saying that perception is everywhere required for finding practicable means, we can say that, to reach such a means, the agent must have a *de agendo* representation of *that* means relative to some target and recognize it as a "*this* that I can do." In the case of fully covert mental action in the absence of an external cue, such *de agendo* representations will have as their targets the vehicles that *would* be appropriate to construct and use in the relevant context. And such representations will come about from the potentiation of the discursive skills or habits whose manifestation *would* result in such construction and use. More on this later.

means of doing *A*]: the possibility it represents as really practicable is one [her] taking of which requires deploying knowledge of this very fact" (2022: 272). Small later identifies the type of fact that an agent in this position believes or knows as the type involved in practical reasoning proper (2022: 275). I take it that the agent's belief or knowledge in this context roughly corresponds to clauses (*b1-2*) of practical want. Namely, such belief or knowledge consists in her taking it (*b1*) that her movement or action is of use towards closing the distance between herself and her end and (*b2*) that the relevant distance is closable by such movement or action. (8) Adding to Ford's (2016), Small argues that, in addition to obstacles and opportunities, the specificity sufficient for an action identified as practicable to actually be practicable depends not just on context but on the agent's skill as well (2022: 271-272). In fact, as an agent becomes more skilled, she becomes more able to implement objectively more and less specific means to her ends.

(9) Following Anscombe, Small identifies three types of fact at issue for practical reasoning. These types are as follows: "type-(i)" facts are those of the form "*one* can do *A* by means of doing *B*", i.e., "doing *A* 'is' a means of doing *B*." "Type-(ii)" facts are those of the form "*I can* do *B* by means of doing *A*", i.e., the now familiar "doing *A* → doing *B*." "Type-(iii)" facts are worldly facts "concerning, e.g., the properties of figures or the location of bakeries" that "underwrite facts of types (i) and (ii)" (2022: 274; original emphasis throughout). (10) In classifying type-(iii) facts as he does, Small allows that they can take any number of forms "e.g. if *p* were the case, *q* would be the case; if an event of type *E* happens, an event of type *F* will happen; *F*s are *G*s; *Fa*, etc." (2022: 274). That is, type-(iii) facts are, for Small, analogous to the objects of speculative knowledge. Thus, they are anything that the agent can theoretically reason about. (11) Small claims that facts of type-(i) depend on there being or having been at least one person for whom the fact in question is or was known as a type-(ii) fact (2022: 275).[47]

Finally (12) Small begins to depart from Anscombe (2005) at this point. Anscombe claims that facts of the form "if plants are fed with certain substances, there will be spectacular plant growth," "if these substances are in the soil, the plants will be fed with them," etc. can feature in practical reasoning, in reasoning towards "a hypothesis proposed for investigation," and in the theoretical reasoning that is familiar to much of philosophy (126-127). That is, there is in Anscombe's thought the idea of, as Small puts it, "a logical kernel that can be used in [at least] two ways, but which is neutral between those two uses" (2022: 275). This kernel would consist of

[47] I take it that Small means to thereby distinguish facts about what one *could do* in principle or in such-and-such circumstances from type-(i) facts. The former type of fact presupposes that the agent does not already have means sufficient to doing whatever it is she *could* do. These are type-(iii) facts. By contrast, type-(i) facts are about someone—possibly oneself—having a sufficient means to doing *B* in virtue of their ability to do *A*.

type-(iii) facts. By contrast, he claims that, given the form of practical reason(ing) as he has developed it, such facts "can be used in theoretical reasoning, but have no place in practical reasoning" (2022: 275). That is, if practical reason(ing) just is connecting series of intentional action concepts to each other with the means-connective, starting from some end which is not practicable to an action that is, then there is no room within such reasoning for facts other than "type-(ii) facts (and their type-(i) counterparts)" (2022: 275). Thus, there is no room in such reason(ing) for type-(iii) facts.[48]

The twelve ways in which Small elaborates the teleological practicality of practical reason give us a nearly complete conception of the capacity to calculate means to ends. This capacity is, again, nothing other than practical reason. So, on the powers framework, practical reason consists in domain-specific discursive sub-skills to order (1) intentional action concepts that (usually) the agent already possesses (2) through the means-connective (4) from some possibly generic or general end to (3) a means sufficiently specific to the agent that involves fully particularized targets (6, 8) that she is skilled enough to immediately instantiate and (5, 7) that she perceives (or non-perceptually recognizes) as a "*this* that I can do." Practical reason thereby represents (9-12) type-(ii) facts having the form "doing *A* (is-à-vis *this*) → doing *B*."

With practical want, we get a more complete conception. The *form* of the exercise of practical reason is captured by Small's formula for stringing intentional action concepts together from any arbitrary intentional action concept that the agent can achieve to a practicable one through the means-connective (i.e., doing *Y* → doing *Z*, doing *X* → doing *Y*, ... doing *A* vis-à-vis *this* → doing *Z*). Practical want gives that form its *matter*. That is, practical want drives practical reason to fix on practicable means both in actual deliberation and in practical reason's manifestation in action by virtue of the agent's taking some end as falling under some good, as at a distance, and as achievable through implementation of the relevant means. Yet, practical want also forms a unity with practical reason because, in practically wanting something, the agent is thereby (already) engaged in movement or action that (she takes) is of use towards such achievement and takes it that what she is aiming at is achievable by such movement or action. Practical reason proper, then, includes practical want. So far, I have been assuming that the form of practical reason proper as Small elaborates it is coincident with its cogency. But in intelligible but infelicitous exercises of practical reason, the agent can be wrong that she can do *B* by doing *A*, that doing *A* is a sufficiently specific means relative to her skill, or that she perceives (or non-perceptually

[48] I discuss this departure from Anscombe in more detail in Appendix B.

recognizes) a practicable means of doing *B* in perceiving what she takes to be a "*this* that I can do." It is to accounting for the cogent exercise of practical reason driven by practical want that I now turn.

1.3.4 Cogent Practical Cognizing and Practical Knowledge

With Springle's naturalistic, non-reductive, action-first account of intentionality on the table, the intentionality of intentional action concepts and, thus, the intentionality of practical knowledge is in view. On my elaboration of the account, practical knowledge represents appropriate means-end orders as appropriate, generalizable solutions. These are solutions whose application *would* satisfy some need in light of the agent's generating the relevant interaction with some target such that *anyone* in possession of the relevant solution system could in principle satisfy that need in light of generating that type of interaction with that target. Exercise of that practical knowledge represents such orders as solutions that the agent is applying or is forming in deliberation so as to apply. That is, such exercise is the activation of the relevant part(s) of the network of her skills or incorporating habits and thereby represents the order to be realized as "I am doing *B* by doing *A* (vis-à-vis *this*)" (or the generation of intentional action acorns for doing *B* by means of doing *A*, where "*this*" corresponds to the proximal target(s)). And with practical want and the teleological conception of practical reason on the table, the practicality of practical knowledge is in view. Other things being equal and assuming no change of mind occurs, practical want either drives search for practicable means to some end through deliberation (itself a mental action) or drives what is effectively affordance-based identification and subsequent[49] implementation of those practicable means. In the latter case, the performance of the action itself instantiates the means-end order expressive of the agent's practical-reason-*cum*- practical-knowledge-(*cum*-intentional-action). I take it that practical reason that is driven by practical want and that takes as its content intentional action concepts which the agent possesses to altogether constitute her practical cognition. In cases of both effective deliberation and non-deliberative knowledgeable performance, I will say that the agent's practical cognition is cogent and achieves the status of an exercise of practical knowledge. What remains, then, is to explain how practical cognition achieves cogency. What makes practical knowledge *knowledge* and not mere belief, opinion, wanting, or even brute desire?

The short answer is that practical cognizing achieves cogency when it represents as a type-(ii) fact a type-(ii) *fact*. That is, cogency in practical cognizing is a function of the agent's representing and therewith recognizing as an order that she can instantiate "an order which is [*really*] there whenever actions [of the

[49] This should not be read temporally but, rather, logically.

relevant type] are done with intentions" (Anscombe 1957/2000: §42). So, the powers framework holds what Schwenkler (2019) calls a "*factualist* understanding" of practical knowledge (99; original emphasis). The possession of practical knowledge is knowledge precisely because it represents as a way of doing something a way that the agent can actually do that thing. In having practical knowledge, the agent is in a position to discursively represent that "I/One can do *B* by doing *A*."[50] And the (complete) exercise of practical knowledge is knowledge precisely because it represents in a way that *makes* a worldly state of affairs the case. That worldly state of affairs comes about by the agent's taking what she recognizes as actually practicable means to her end and actually achieving that end through those means. Anscombe, in discussing her telling someone that she is opening a window, puts it nicely: "I have called such a statement knowledge all along; and precisely because in such a case what I say is true—I do open the window" (1957/2000: §28).[51]

The key to the powers framework's factualist understanding of practical knowledge is the claim that intentional actions are developing processes which have an *a priori* teleological unity consisting in moments of activation of intentional action concepts, ongoing performance, and satisfaction of some need. This *a priori* teleological unity is ultimately grounded in the agent's possessing internally calculatively structured skills and incorporating habits. Among the upshots of this claim is that the outright failure of practical cognition requires more than that the agent is at present not acting or is failing to act as she intends (Falvey 2000; Thompson 2008; Small 2012). In the latter case, such failure is (usually) only *prima facie*. The intelligence of cogent practical cognizing, then, consists in its robust flexibility and non-accidental success regarding the agent's taking as objects future contingents that she can affect, ordering intentional action concepts from those future contingents to practicable means, and implementing those means in intentional action.

That robust flexibility manifests in finessing or correction of ongoing performance and respecification of the agent's intention in light of *prima facie* failures of practical cognition (Small 2012: 158). Such failures consist in the agent's failing to recognize something that is actually affectable by her action, failing to recognize actually practicable means to those ends, or failing to have speculative knowledge of some fact upon which the success of her action depends.[52] In all three cases, if the agent were to express her intention, her statement would

[50] I will have more to say about recognition in the following chapter.

[51] Note that she does not say "because what I say is expressive of a corresponding propositional attitude whose truth-conditional content—that I am opening the window—is made true by the fact that I am opening the window."

[52] To ward off common concerns about the factualist understanding, I address *prima facie* failures of practical cognition and mistakes in performance in more detail in Appendix C, Davidson's carbon copier case in Appendix D, and Bratman's argument against the Simple-View in Appendix E.

"fall to the ground" (Anscombe 1957/2000: §32; *cf.* Small 2012: 146*ff.*). Such an expression falls to the ground whenever the agent is at the moment not satisfying her intention by whatever means she takes herself to be implementing. But the agent can pick her statement back up by respecifying her intention and going on to execute that respecified intention or by fixing on means by which she can actually affect what she is aiming to and implement those means (Small 2012: 147, 157, 202).

All that is left, then, are the infamous "mistakes in performance" or what Anscombe calls "the *direct* falsification" of an expression of intention (1957/2000: §32; original emphasis). In such cases, there is no failure of practical cognition. As Anscombe, echoing the author of *Magna Moralia*, puts it: "the mistake is not one of judgment but of performance" (§32). This means, among other things, that the agent's practical cognition is cogent when she makes a mistake in performance. And if she has opportunity to correct for such mistakes, her cogent practical cognizing achieves the status of an exercise of practical knowledge in her correcting the mistake. That is, given our finitude as agents, such mistakes are accidental and unlucky performances or aspects thereof that the agent can usually correct. And it is thanks to the *a priori* teleological unity of action that this is so. That is, mistakes in performance are defects in the development from activation to ongoing performance or from ongoing performance to success. Moreover, it is thanks to this unity that in one-off cases where such mistakes occur, the agent can nonetheless say *truly* of what was happening that she was *doing* what she nevertheless never did (§23). After all, someone who survived being hit by a bus can start their story with "So, I was walking across the street…"

So, cogent practical cognizing must tolerate and be sensitive to *prima facie* failures as well as mistakes in performance. Cogent practical cognizing must also in some way draw on speculative knowledge (or, at minimum, correct opinion) of at least two subsets of type-(iii) facts. First, the agent should know type-(iii) facts that underwrite her being in a position to affect her ends. Second, she should know type-(iii) facts that underwrite her being in a position to take actually practicable means to those ends, especially in cases where success is not a matter of course. The successful exercise of cogent practical cognizing is the exercise of practical knowledge. That is, cogent practical cognizing in the absence of mistakes of performance is the non-accidental making true of an agent's doing *B* by means of doing *A* vis-à-vis *this*. This leaves out some details of how practical reason avoids settling on impracticable means, how practical want avoids settling on unattainable ends, and how performance avoids failing to live up to the agent's otherwise cogent practical cognizing. The agent's ability to (re)structure her own agency in recognition of *prima facie* failures and mistakes in performance as well

as in forming intentions for the future fills part of this gap.[53] The contributions that speculative knowledge makes to practical knowledge fills in the rest.

Starting with structuring agency for future action, according to Small (2012) "the true ground and meaning of [the agent's practical knowledge], properly understood, lies in the (i) *calculative* and (ii) *temporal* structure of intentional action" (135; original emphasis). We have already seen what the calculative structure of intentional action consists in, namely, the ordering of intentional action concepts from some end to practicable means, or, in other words, her doing *A* vis-à-vis *this* → doing *B*. The temporal structure of intentional action is important for understanding future action and action that extends over a large period of time. Yet, such structure is part and parcel with the calculative structure. After all, the development of an intentional action from activation of the relevant intentional action concepts through to the complete instantiation of the corresponding means-end order in successful performance takes place and unfolds in time. Both are dimensions of the moments constituting the *a priori* teleological unity of intentional action. Part of the agent's exercise of practical knowledge, then, consists in knowing *in prospect* "that it will be no accident if [she] ends up having intentionally done what [she] is doing intentionally (or intends to do)" (136). Small goes on to suggest that such knowledge in prospect depends on the agent's having the relevant procedural knowledge-how. That is, such knowledge in prospect depends on the agent's knowing she can do *B* by means of doing *A*. Such procedural knowledge-how is in turn grounded in the internal calculative structure of her skills and incorporating habits (136-137). Indeed, Small ultimately claims that "intentional actions, through their calculative articulation, underwrite their own temporal unity: *the calculative structure of an intentional action constitutes its temporal structure*" (189; original emphasis; for a more thorough defense of this claim, see Appendix C). So, an at least proficient driver has practical knowledge in prospect, say, that he is *going to* pick up his kids from school by virtue of the calculative structure instantiated in his heading towards his car, getting in, turning on the ignition, and driving in the direction of his kids' school. Whether there is a "break" between the driver's heading to his car to drive to his kids' school and his picking them up depends on the degree to which the latter is a matter of course for him (Anscombe 1957/2000: §34). Nonetheless, even if picking up his kids is not a matter of course, what makes it the case that he has practical knowledge that this is what he is *going to do* still just is his cogent practical cognizing the means to that end.

[53] I discuss (re)structuring agency in light of *prima facie* failures and mistakes in performance in detail in Appendix C.

It is worth re-emphasizing at this point that the powers framework does not conceive of intentions as mental events that exist independently of their objects and stand to them as isolable truth-conditional representations that (somehow) non-deviantly bring them about. Intentions are not a required "extra feature" that structure or guide action by (efficiently) causing it (*cf.* Anscombe 1957/2000: §48). Rather, they are the activation of the relevant intentional action concepts themselves in the active synthesis of deliberation or in the immediate instantiation of what will become the agent's intentional action. That is, intentions represent intentional actions by virtue of being a phase of the latter's coming about and stand for (not *in* for) them as appropriate for satisfying the need that those actions function to satisfy.

In deliberation, the goal is, in effect, to form an intention to perform some other action. In forming that intention, other things being equal and assuming the agent does not change her mind, she begins acting towards that goal *then and there*. For example, suppose one morning an agent realizes that she is nearly out of food. She deliberates and determines that she can get food by purchasing it, can purchase it by going to the grocery store, and can go to the grocery store by driving there after work. Her judgment that this is what she *will do* immediately structures her action going forward. Her outward behavior from then until the end of her workday might be completely indistinguishable from such behavior on any other day. Nonetheless, it is true *all along* that she is *going to* go get food by buying it at the grocery store, that she is *going to* buy it at the grocery store by going there after work, and that she is *going to* get to the grocery store by driving. Reductive causalism attempts to explain this by appeal to her forming a truth-conditional propositional directive attitude that stands apart from what she will do and its success in later (non-deviantly) causing or guiding her action when the time comes. By contrast, the powers framework insists that, from the point of making the judgment, the agent is literally already acting towards this goal.

That is, on the powers framework as I am presenting it, the agent's, say, *going to* go get groceries after work is something that she is *doing* the whole time and consists in the partial activation of the relevant intentional action concepts appropriately ordered into the relevant means and end. Such activation is only partial for at least one of two reasons. First, it might be that the relevant target of the agent's *de agendo* representation for going to the grocery store does not yet obtain, say, anything that indicates that the workday is over. Second, it might be because the agent's solution system for going to the grocery store is suppressed, say, by her solution systems for going to work and working. Or activation might only be partial due to both factors. All that is required for activation sufficient for initiating performance is the impact on the agent of the obtaining of the relevant target-fact, say, the clock's reading "5:00PM," or the removal of suppression coming from her work-

related solution systems (or both). In both kinds of case, the calculative structure is present in the pattern of activation of the network(s) of intentional action concepts constituting her skills or incorporating habits *all along*. So, the agent's intention to go get groceries after work is there *all along*. The presence of the intention does not require an additional truth-conditional propositional attitude that tends towards structuring her action from the outside. It is the agent herself in the capacity of an integrated complex of skills and incorporating habits that does the structuring. That is, it is the agent *qua* agent that structures her action in modifying her solution systems.[54]

So, beyond cogent practical cognizing considered by itself, what allows the agent to latch onto achievable ends and practicable means such that she knows that it will be no accident if she ends up having done what she is intentionally doing or intends to do? Part of the answer lies in the role that speculative knowledge plays in buttressing cogent practical cognizing.[55] As was noted, in many cases, perception guides recognition of a practicable means. This does not mean that the agent knows *that* she is doing *A* (vis-à-vis *this*) on the basis of perceiving that she is doing *A* (vis-à-vis *this*). Rather, her perception is of a target and represents a practicable means involving that target as a "*this* that I can do" or as "I can do *A* (vis-à-vis *this*)" where doing *this* is an exercise of one of her incorporating habits. Moreover, perception and inference concerning type-(iii) facts

[54] That said, I grant that—especially with actions that extend over a period of time—the agent can lose track of what she is doing. Human agents are fallible and one way our fallibility manifests is in practical failures obstructing or impeding the full realization of the structure we instantiate in the world through whatever action we set upon. Sometimes this takes the form of *prima facie* failures of practical cognition. Other times it takes the form of degradation in structuring our actions in accordance with the calculations that we have settled on. When an agent absentmindedly takes her route home when she intends to go to the grocery store, this mistake in performance reflects such degradation. Prospective memory, reminding ourselves of what we are doing and of what we are aiming at, and keeping track of our intentions are all forms of actively maintaining that structure in the face of distraction, competing desires and intentions, and a world that is not always so accommodating. Similarly, some of our reliance on speculative knowledge to aid calculation and execution are likewise ways of maintaining that structure. Acknowledging this does not require positing a separate mental state with truth-conditional contents that persists from its formation throughout the duration of whatever else the agent is doing or throughout performance of the relevant action. All it requires is that the relevant part(s) of the relevant dispositions be structured accordingly and that their being (partially) activated does not automatically equate to (outward) performance of the relevant action. I also grant that in cases of actions that are highly abstract and that extend over large periods of time like, say, getting a PhD, the agent at no point has a single solution system whose (partial) activation is maintained throughout the course of performing the action. Rather, such an action requires the activation and application of a host of interrelated solution systems that are brought together under (at least the potentiation of) a discursive representation of the relevant plan(s).

[55] The other part of the answer lies in distinguishing between cases where an agent intends to do something that she is not yet skilled at and cases where she merely "aspires" to do such a thing (Small 2012: 167*ff*.) I discuss such cases in Appendices D and E.

relevant to the success of an action can underwrite the identification of future contingents affectable by the agent as well as (practicable) means.

For instance, an agent might infer from seeing what time it is that he *can* pick up his kids from school in the next hour and, so, better start making his way to the car. The former is the relevant future contingent and the latter is the corresponding practicable means. His inference does not provide evidence or demonstrate to him that he *is* now going to pick up his kids from school. Rather, his looking at the clock and knowing the scheduled times for pick-up allow him to draw an inference to the relevant type-(iii) facts which underwrite what are for him the type-(ii) facts that *he can now* go pick his kids up from school by driving his car there by making his way to his car, getting in, and turning the ignition vis-à-vis *this*. Like the grocery shopper, our inattentive father's inference functions to structure his agency. Yet, unlike the grocery shopper, his action of going to go pick up his kids from school is not something he is doing *all along* as he works. Rather, he is structuring his agency on the fly: the network of intentional action concepts for getting to the school by car are not even partially activated until he sees the time and infers that his kids will soon be available for pick-up.

On Springle's account, such perception and inference involve activation of multiple solution systems (or the generation of multiple practical-inferentially related intentional action acorns) culminating in the modification and application of the agent's solution system for, e.g., returning his kids home. Perception of the clock affords recognition that it will soon be time to pick up the kids. Such recognition is at minimum the activation of a set of interrelated discursive solution systems connected by having the function to discursively represent the same target (more on this in chapter 2). In this case, the target in question is the fact that it is almost time to pick up the kids. Either the agent actually constructs and uses a discursive representation, say, in muttering to himself "Ah, better go pick up the kids," "Ah, it is time to pick up the kids," etc., which then generates *de agendo* representations, say, of his desk as get-uppable-from, his office as leavable-from, his car as get-intoable and drivable, and so on. Or he draws on the discursive contents that *would* be applied in that discursive representation in generating those *de agendo* representations directly. Both are transitions in thought by virtue of being transitions from *de agendo* representations to (at least) activation of discursive skills to further *de agendo* representations (more on this in chapter 2). The latter case involves a fully covert transition in thought because the agent does not actually apply the discursive solution system(s) in question. He does not actually say (or write, draw, map, diagram, etc.) anything. He merely trades on activation of the relevant discursive solution system(s) in generating the relevant opportunities for action in anticipation of a discursive representation that will never actually come. In both cases, his generating the relevant opportunities modifies his

solution system for returning his kids home by drawing on his intentional action concepts for his most immediately practicable means in the relevant context, bringing them under to the network of intentional action concepts he possesses for picking his kids up. In this case, the former are his intentional action concepts for getting up from his desk, leaving his office, walking to where he parked his car, and so on. His exploitation of these opportunities, in turn, just is his applying those sub-solutions as part of applying his solution system for returning his kids home.

What's more, for actions involving multiple isolated steps, the agent will "often make use of [her] senses," gathering evidence that she has completed one step before moving on to the next (*cf.* Anscombe 1957/2000: §48). For instance, hearing the kettle whistle lets the agent know (read: recognize) that it is time to take it off the stove. Yet, it is not by hearing that whistle that the agent has practical knowledge of boiling the water, say, to make tea. This is because her practical knowledge of boiling the water is grounded in the calculation of practicable means towards the end of making tea. Even if she were to check that the tea was sufficiently well-steeped before going on to drink it, this would not of itself show that her practical knowledge of making tea was based in perception *that* she is making tea. Rather, her checking and perceiving that the tea is sufficiently well-steeped provides recognition of the obtaining of a target-fact, namely, that the tea is ready to drink. Such knowledge is an aid to the success of action and to the performance of further actions. [56] In general, perception and inference support successful practical cognizing by forming a bulwark both for preventing *prima facie* failures and against those failures going unaddressed.

Neither perception (by itself) nor inference, then, take the place of cogent practical cognizing. When the agent sees or infers that what she thought was a future contingent affectable by her is not at this time so affectable, she *then* engages in deliberation or intentional action with the goal of turning this currently unsolvable problem into one that she can solve. When she sees or infers that what she thought were practicable means are in fact not practicable, she *then* engages in deliberation or intentional action with the goal of fixing on

[56] As Anscombe (1957/2000) puts it:

> but isn't the role of all our observation-knowledge in knowing what we are doing like the role of the eyes in producing successful writing? That is to say, once given that we have knowledge or opinion about the matter in which we perform intentional actions, our observation is merely an aid, as the eyes are an aid in writing. Someone without eyes may go on writing with a pen that has no more ink in it; or may not realise he is going over the edge of the paper on to the table or overwriting lines already written; here is where the eyes are useful; but the essential thing he does, namely to write such-and-such, is done without the eyes. So without the eyes he knows what he writes; but the eyes help to assure him that what he writes actually gets legibly written (§29).

actually practicable means. When she sees or infers that some type-(iii) fact on which the success of her action depends does not obtain, she *then* engages in deliberation or intentional action either with the goal of making it obtain through realizing a corresponding type-(ii) fact or with the goal of achieving her end in a way that does not rely on that type-(iii) fact. In any of these cases, of course, an alternative is to abandon the relevant bit of practical cognizing, giving up the corresponding intention. But the point is that perception and inference constitutive of speculative knowledge functions to aid successful action in myriad ways. And when an agent sees or infers that all is well, that she is pulling her action off, her practical knowledge is already present. That is, her practical knowledge that she is doing *B* by doing *A* vis-à-vis *this* is not based on her perception or inference even in this case. If this is right then, again, "we really can speak of two knowledges" of one object: the one practical and constitutive of that object, the other speculative and at most *an aid* in the production of that or further such objects (§48). What perception and inference do for the agent in all of these cases is inform her of type-(iii) facts relevant to the success of her performance and of what she can go on to do. She does not gain practical knowledge of doing something, why she is doing it, or how she is doing it through either perception or inference. Rather, she has and exercises practical knowledge of doing things by virtue of her deliberation or acting *sub specie agentis*, as it were.

The exercise of cogent practical cognition—with appropriate support from speculative knowledge—includes the agent's doing what she can to compensate for her fallibility. Cogent practical cognizing is cogent in part by ultimately fixing on ends that are actually attainable by the agent and means that she can actually implement, where the relevant practicable means ultimately fall within her repertoire of incorporating habits. Likewise, such cognizing is cogent and robustly flexible in its activity by virtue of admitting adjustment and refinement of her calculation as she acts. It is only by taking cogent practical cognizing to be some (occurrent) truth-conditional propositional attitude somehow entailing, containing, or fashioned after propositional attitudes like belief that there is any temptation to move away from thinking of practical cognition as being activation of the relevant network(s) of intentional action concepts themselves and, thus, as being of the essence of intentional action itself.

On the powers framework, then, in order for an agent to be intentionally doing something, she must have practical knowledge of her doing it (as well as why and how she does it). Her cogent practical cognizing consists in her fixing on actual future contingents affectable by her action and her calculating actually practicable means to those ends. In Springle's framework, this amounts to her fixing on actual solutions for actual problems and the targets that those solutions involve, on the one hand, and on actual sub-solutions

whose application would constitute the application of those solutions on the other. Such cognizing is buttressed on one side by its own robust flexibility in addressing *prima facie* failures and mistakes in performance as well as in its role in structuring one's own agency. It is buttressed on the other side by the agent's speculative knowledge of type-(iii) facts that are relevant to the success of her action. And such cognizing is essentially active: it just is her deliberating with a view to immediately instantiating the relevant embodied means-end order or is her simply immediately instantiating that order, thereby non-accidentally making true her doing *B* by doing *A* vis-a-vis *this*.[57] This is the power framework's positive factualist account of practical knowledge as cogent practical cognizing. I now turn to fleshing out the relationship between the capacity for practical knowledge as cogent practical cognizing and other skills and habits.

1.4 Tying the Threads Together

I said that cogent practical cognition, i.e., practical knowledge, is a capacity for ordering means to ends but that it is not one that exists independently of the skills and habits whose manifestations in performance are the exercises of such knowledge. I also said that skill and habit can be characterized by intentional action concepts and the calculative structure embodied in their means-end coherence as well as in their practical-inferential coherence. Finally, in drawing on Springle's account of intentionality, I said that an agent in a position to act represents some order as appropriate for satisfying some need by virtue of the activation of those intentional action concepts (or the generation of intentional action acorns for those actions) either in the active synthesis of deliberation or in the dual activation of *de agendo* representations and discursive activity that develops into the immediate instantiation of an intentional action. In this section, I tie these claims together with the metaphysics of skill and habit presented in Section 1.2. To this end, the powers framework holds the following two claims. First, practical cognizing is not something over and above an agent's skills and habits. Second, acquisition of the latter involves coming to know how to order the intentional action concepts specific to their domains into

[57] Compare:

> If we hear of a 'capacity,' it is reasonable to ask what constitutes an exercise of it. [...] In the case of practical knowledge the exercise of the capacity *is nothing but the doing or supervising of the operations of which a man has practical knowledge*; but this is not *just* the coming about of certain effects [...] for what he effects is formally characterised as subject to our question 'Why?' whose application displays the A—D order which we discovered. [...] 'Intentional action' always presupposes what might be called 'knowing one's way about' the matters described in the description under which an action can be called intentional, and *this knowledge is exercised in the action* and is practical knowledge (Anscombe, 1957/2000: §48; emphasis added).

appropriate means-end calculations. The mutually reinforcing interdependence of habit and skill is therefore further illuminated with the introduction of practical cognizing and Springle's account of intentionality: the calculative capacity is brought into service in the rational transmission of skill and is at work throughout the life cycle of a skill in the modification and development of ever more complex solution systems. To show how this works, I will be giving an even finer-grained gloss on the life cycle. However, before getting there, I want to say a little more about propositions that apply intentional action concepts and so-called "basic action." It is to this that I now turn.

1.4.1 Practical Propositions, Codified Procedural Knowledge-How, Social Praxes, and Practical Reason

The viability of the powers framework depends in part on finding a place for propositional content. Yet, I have saddled the framework with the task of explaining intelligent agency in the absence of positing (occurrent) truth-conditional propositional directive attitudes whose function is to (non-deviantly) initiate and control or guide behavior in an intelligent way. It is not that the framework denies the existence of such attitudes. After all, it commits to their existence, though in a radically different way than does its causalist rivals. Where its rivals posit these attitudes to explain intelligent agency by modeling them after attitudes like belief, the powers framework understands them as the products of the application of discursive solution systems. These solution systems and their application are essential both to the rational transmission of skill and to the agent's ongoing rational self-shaping. I will claim in chapter 2 that some overt mental actions just are such applications. So, the problem is not that the powers framework does not make room for propositional attitudes or does not hold that there is a relationship between them and intelligent agency. The problem is that I have prohibited the framework from appealing to these attitudes in accounting for the intelligence of each particular intentional action as what (non-deviantly) causes those actions. If such attitudes were required to be the (non-deviant) cause of each particular intentional action then, again, the powers framework would not be a legitimate alternative to its causalist rivals. So, the question is what role propositions play in the power framework's account of human intentional action. The answer, again, is that propositions that apply intentional action concepts facilitate the rational transmission of skill and the rational self-shaping of which human beings are capable. This answer suggests that such propositions are partly definitional of the sorts of social praxes that are necessary for the very existence of the skills that they impart.

Recall that intentional action concepts have the form "doing *B* (by doing *A*)." Propositions applying such concepts, in turn, are discursive representations with the form "I can do *B* by means of doing *A* (vis-a-vis

this)" or "doing *A* (vis-a-vis *this*) → doing *B*." On the powers framework as I am presenting it, these propositions come onto the scene with the formation of a social praxis as codified ways of doing something in some domain. They consist, in the first instance, in publicly available discursive representations encoding procedural knowledge-how of domain-specific solutions to domain-specific problems and represent type-(i) facts. For instance, the aesthetic values at work in the praxis of playing piano make dexterity and fluid movement of the hands and fingers extremely important. This poses a domain-specific problem insofar as some hand positioning and postures threaten to be detrimental to such dexterity and fluidity in play. The solution is identifying type-(i) facts like "One can minimize stiffness of the fingers and maximize dexterity by positioning one's hands thus-and-so." And, given the dependence of type-(i) facts on at least one agent's representing or (at some point) recognizing a corresponding type-(ii) fact, this means that some pianist must have at some point achieved at least possession of practical knowledge of, say, "I can minimize stiffness of the fingers and maximize dexterity by positioning my hands thus-and-so (vis-a-vis *this*)" prior to codifying that knowledge in the form of some discursive representation of the corresponding type-(i) fact in speech, writing, drawing, etc. so as to teach others. In that moment, this imagined pianist, on the one hand, activates the relevant intentional action concepts and their embodied means-end order and, on the other, activates corresponding intentional action concepts for the construction and use of a discursive representation that applies the former intentional action concepts. That is, activation sufficient for performance of the latter intentional action concepts develops into the agent's constructing and using a discursive representation with the content "One/I can minimize stiffness in the fingers and maximize dexterity by positioning one's/my hands thus-and-so." And if the imagined pianist is truly making a discovery about how to play piano in this moment, she is also structuring her own agency by ordering the relevant intentional action concepts into the relevant means and end. More on this later.

All procedural knowledge-how whose application in discursive activity or whose being drawn on in acting intentionally is worthy of being called an exercise of intelligent agency starts out in this in principle publicly available form. Otherwise, the existence of social praxes would not be necessary for the acquisition of skills or the category of skill would be tightly circumscribed to expert skill. For whatever procedural knowledge-how started out as private, internally available contents of corresponding propositional attitudes would in all likelihood concern the performance of actions constituting the elementary atoms of skills. An agent could not acquire these attitudes by way of the relevant social praxis, since such praxes are by definition public and interpersonal. And if she could string these atoms together into the relevant complex skillful actions without the aid of the praxis, then the praxis would not be necessary for acquiring the relevant skills. All the agent would

need is sufficient habituation stringing the relevant atoms together. If the social praxis and, thereby, guidance by some expert were required to string them together appropriately then, even so, the agent would only require the praxis for acquiring those skills she could not master just by having the relevant atoms. In all likelihood, such skills would go beyond those constituting the performance of everyday intentional actions. They would be the so-called "expert skills."

Assuming, then, that the existence of social praxes are necessary for the acquisition of skills in general and that the category of skill applies just as much to everyday skills of the sort we recognize as capacities for mundane intentional actions as it does to expert skills, propositions like "One/I can tie one's/my shoes by arranging the laces thus-and-so" and like "One/I can play *Wrong Note* by playing the piano thus-and-so by..." must begin their lives as in principle publicly available domain specific procedural knowledge-how. Call propositions with such contents "practical propositions." The foregoing notwithstanding, the contents of practical propositions in the first instance pick out possible interactions underwritten by networks of internally calculatively structured parts of skills or incorporating habits as those networks are embodied in individuals. And when applied in the expression of intention or in action description, explanation, or assessment, their content picks out the order that is actually there when actions of the relevant type are performed with intentions.[58] Again, their content is coincident with that of the intentional action concepts activated in the first moment of the *a priori* teleological unity of the action that they represent.

On the powers framework, intentional action concepts are at work in the exercise of procedural knowledge-how. In the expert performance, such concepts are drawn on in manifestation of their skill. However, such concepts are also applied in discursive activities for constructing and using the relevant practical propositions. The master pianist's playing puts her in a position to express her intention and to describe, explain, or assess what she is doing because she has and is exercising practical knowledge of what she is doing,

[58] This leaves two issues unresolved: how do social praxes come onto the scene in the first place and how can type-(i) facts depend on type-(ii) facts if all propositionally articulated procedural knowledge starts its life out as in principle publicly available contents that represent type-(i) facts? Resolving both issues requires giving an account of animal agency and the coming about of praxes that allows for continuity between animal agency and rational agency without, on the one hand, attributing rational agency to animals and, on the other, portraying rational agency in human beings as an additional isolable capacity that merely sits atop animal agency. While I do not have the space for it here, the resolution is, I think, essentially two-fold. First, we ought to ascribe volitional dispositions to animals that can be transmitted arationally through interaction and child-rearing. Second, we ought to hold that social praxes that engender skills come onto the scene as responses to communal problems born of our deep mutual interdependence on one another for thriving. The first instances of intelligent agency were at once expressions of type-(ii) and type-(i) facts, establishing the formation of the first social praxes. The ratchet of cultural evolution, then, would start with many cases like that of our imagined pianist.

why she is doing it, and how. That is, the activation and drawing on of the intentional action concepts for playing by the master in acting in turn activate corresponding intentional action concepts for constructing and using corresponding discursive representations. Yet, the intentional action concepts for playing have their calculative structure in the first place thanks to application of those concepts in the construction and use of corresponding discursive representations. In the first instance, such discursive representations, as constructed and used by her mentors in instruction and guided practice, provide a structure for the agent's action and norms of success that the agent approximates. After acquisition of the skill, her (at minimum) potentiating the construction and use of such discursive representations in structuring her own agency provides a structure and norms of success that the agent aims to realize and satisfy, respectively, on the relevant occasion. The mutual interdependence of intentional action concepts for playing and intentional action concepts for discursively representing play suggest that the latter are not components of a discursive skill that floats free of the skill for playing. Rather, such discursive skills are determinates of the shape *in concreto* of the skill for playing. And the master's ability to describe, explain, or assess the actions specific to the domain of playing piano in the absence of playing herself draw on those very same determinates and, thus, depend on her having the corresponding ability to play.

Each skill, then, has its own "language" with its own range of idiolects that must be learned as the skill is learned and one that an agent possessing the skill adds to when she adds to the edifice of the corresponding social praxis. The powers framework denies that there is a general linguistic capacity for representing actions that exists over and above the integrated complex of skills and habits constituting the agent as agent. For language needs a foothold in domains and their objects in order to convey or appropriately transform meaningful statements. That said, what makes the discursive determinates of multiple skills count as capacities for discursive representation (including language) across those skills is their bringing to bear fully generalizable calculative and practical-inferential structure onto the shape *in concreto* of those skills. It is that generalizability that allows the exercise of these discursive determinates to give ampliative application of determinate skills and subsumed habits from the shape *in concreto* of their determinable skill to the integration of those dispositions in the formation of the shape *in concreto* of a novel target skill.[59]

[59] In fact, it's entirely consistent with the powers framework that facility with enough medium-sized dry goods brings with it linguistic and reasoning determinates that abstract predicate-argument structure, quantification, and so on from sufficient skillful interaction with these objects. This does not obfuscate the fact that developing the linguistic and reasoning determinates of other skills does not thereby come for free despite each domain of skill arguably exhibiting the same predicate-argument, quantificational, etc. structure among the objects of its domain. Thus, the powers framework

In the novice, the work of domain-specific intentional action concepts is primarily done in the manifestation of determinate skills or subsumed habits that are drawn on in the performance of actions within the domain of a novel target skill and which are guided by an expert manifesting discursive and pedagogical determinates of that target skill. It is by virtue of such guidance that the expert structures the drawing on of the relevant intentional action concepts such that they instantiate the appropriate means-end order in the novice's action. However, these concepts are also applied by the novice himself through the additional activation of discursive determinates of some skill(s) he already possesses for expressing his intentions or describing, explaining, or assessing his or others' actions within the domain of the relevant home skill(s). For instance, the novice pianist is not yet able to fully or accurately express his intention in the pianist's language. Nor is he yet able to fully or accurately describe, explain, or assess in the pianist's language what he is doing as he is learning to play. But he is able to do either using one of his other "tongues," say, the typist's language or at least the languages attaching to the sorts of everyday bodily actions that form parts of a variety of skills that children learn in the first few years of life. That is, the novice can find rough equivalences between say "placing one's/my hands thus-and-so" with a view to playing piano and "placing one's/my hands thus-and-so" with a view to, say, typing on a keyboard. Both his drawing on those intentional action concepts in learning to play piano and his application of them in expressing his intention or describing, explaining, or assessing the relevant action are ampliative in their finding appropriate use in the target domain of playing piano. As the novice approaches proficiency, the skills that will become determinates of the target skill's shape *in concreto* and the habits that will be subsumed into that shape are structured by his teacher with a view to his acquiring the relevant action concepts. His acquiring them at once brings with it acquisition of the novel skill and of corresponding discursive determinates for expressing his intention and describing, explaining, or assessing actions within the domain of that novel skill.[60]

---

need not deny an Evans-style generality constraint. An agent can come to recognize the predicate-argument structure in one domain and then need only grasp the concepts particular to another to start applying that structure to the latter. I take it that such generalization is part of how novices latch onto intentional action concepts in new domains as part of integration.

[60] This way of describing the dynamic between novice and expert ignores instances of self-teaching and the discovery of new social praxes. The former can be seen as an extension of this dynamic. The novice treats as expert teacher codified procedural knowledge-how itself. The latter requires its own treatment as part of a larger account of rational transformation. Similarly, this way of describing the dynamic touches directly on the contemporary know-how debate in philosophy of action. The debate is characterized by two camps: intellectualism and anti-intellectualism. Intellectualists hold that procedural knowledge, at bottom, consists in propositions that the agent accesses through the relevant propositional attitudes and applies in action by virtue of these attitudes having a "practical mode of presentation" (for instance, see Stanley and Williamson 2001; Pavese 2015, 2020). Intellectualists are therefore strong allies of the causalist

Ampliative drawing on and application of intentional action concepts from one domain to another points us back to what I call "integration" (Section 1.2.2). Recall that integration occurs when, in coming to acquire the integrating skill, the agent gives the integrated skill or habit ampliative application to the domain of the target skill. So far, I have for the most part been concerned with the acquisition and mastery of individual skills in isolation. But integration suggests that skills beget skills not just in the expert's rationally transmitting her skill to the novice but in the novice's making use of a skill he already possesses in acquisition of new skills as part of such rational transmission. Indeed, integration reveals how skills are interdependent. No rational agent starts out with just one skill. Rather, her achieving the status of rational agent brings with it a host of integrated skills. Integration likewise prepares the agent to acquire novel skills by virtue of her being part of her latching onto fully generalizable structures of means-end orderings across the domains of the battery of skills that she acquires as part of coming into rational agency. And her rational agency consists just in this, namely, the integrated set of skills that she possesses. Her ability to coherently order intentional action concepts and

---

picture. Anti-intellectualists hold that procedural knowledge, at bottom, consists in abilities to perform the relevant action, where this need not involve any propositional contents or attitudes. Anti-intellectualists are therefore strong allies of (neo-)behaviorism. The powers framework claims that procedural knowledge-how starts out as in principle publicly available propositional contents. It thus sounds like the framework is intellectualist. Yet, it also says that an agent's procedural knowledge-how is nothing other than her ability to perform the relevant actions and is both expressed and exercised in such performance. It thus sounds like the framework is anti-intellectualist. So, the powers framework appears to be at odds with itself.

On the contrary, the framework proposes a resolution to the know-how debate by denying that propositional directive attitudes have the role of initiating or controlling the performance of intentional actions. It agrees with intellectualists that procedural knowledge-how worthy of counting as knowledge-how to engage in intelligent action depends on such knowledge being in principle articulable in speech, writing, picturing, maps, etc. It agrees with anti-intellectualists that such procedural knowledge-how as it is instantiated in human beings consists in the relevant abilities and their exercise. The framework holds these two claims together by holding two additional claims. First, the relevant abilities for action just are the agent's possessing the relevant intentional action concepts where their being drawn on just is her manifesting those abilities in intentional actions. Second, in order to have those abilities, the agent must have at some point learned them, where learning them depends on her being brought into a social praxis that, *inter alia*, codifies this procedural knowledge through external publicly available propositional contents in, e.g., speech, writing, pictures, maps, etc. The relationship between the embodied form of propositionally articulable procedural knowledge-how and its codified, articulated form is what explains the oft cited case in the know-how literature of the coach who cannot do but can teach. Intellectualists take the coach to possess the relevant know-how, whereas anti-intellectualists deny that he does. The powers framework would ascribe the coach knowledge-how *to convey* the relevant codified procedural knowledge-how in the form of a corresponding pedagogical skill. That said, the framework would also contend that the quality of what the coach can convey is likely to covary with how close he got to acquiring the target skill himself. Olympic coaches provide a good case here: they cannot engage in whatever sport at an Olympic level but they are competent enough that they can push others to the Olympic level through the exercise of the pedagogical determinate of the shape of their skill. A coach who could not perform the most basic acts in the relevant domain is unlikely to be an adequate pedagogue with respect to that domain. So, the possibility of such a coach in possession of the relevant know-how, while arguably conceivable, does not of itself show that knowledge-how consists in propositional attitudes.

instantiate those orders across a number of related domains constitutes the unity of her rational agency and her exercise of this ability in both deliberation and intentional action is the exercise of that agency.

Again, the powers framework subscribes to a transformative theory of rationality. Intentional action concepts and their possession are not solely the jurisdiction of discursive representation. What makes both skills and habits properly rational (and, so, intelligent) is that, once acquired, they are shot through with a calculative structure that is in principle articulable in discursive representations and that allows for ordering in accordance with in principle fully generalizable principles of practical reason(ing) concerning type-(i) and type-(ii) facts and theoretical reason(ing) concerning type-(iii) facts. It is in virtue of being articulable and of admitting of such ordering that the agent represents her own intentional action concepts in both *de agendo* representations and in her being primed to discursively represent orders that any agent situated as she is could instantiate in action. By contrast, nonhuman animal action, while ordered and arguably intentional (Anscombe 1957/2000), is not shot through with this rational calculative structure.

That nonhuman animal action is not shot through with rationality is arguably evidenced both by an inability on the part of such animals to express knowledge of what they are doing with discursive representations and by the lack of full generalizability of their action concepts across an open-ended range of domains. Animal action concepts are likely tied to their domains in a way that intentional action concepts are not. While intentional action concepts are acquired relative to a domain, they can be drawn on and applied in the acquisition of a seemingly unlimited number of novel skills. The form of rational agency begins to inhere in individual human beings once they've extracted it from the battery of skills they acquired in early development. By contrast, animal action concepts seem to persist and pass away in or in close proximity to the domain in which they were acquired. Animals learn novel skills and, to do so, must give ampliative application to action concepts they already possess in a related domain. So, there is some generalizability of action concepts for animals. Yet, animals do not seem able to extract a fully generalizable structure that is open-ended in its application to the acquisition of novel skills. By contrast, in developing the capacity for practical reason, humans latch onto such a structure, namely, the calculative structure. I contend that such latching on is a product of the initial integration of skills constitutive of an agent's "coming to the age of reason." So, while animals acquire capacities for action whose exercise can be ordered so as to facilitate successful intentional action and while they can acquire novel skills, they do not extract the means-end order *as such* nor are able to generalize it such that they can discursively represent the very action concepts that those abilities embody. They

therefore do not develop the sorts of social praxes that allow for the rational transmission of skill and their skills cannot be fully unified through integration.

Similarly, in line with its commitment to a transformative theory of rationality, the powers framework rejects the claim that there is a capacity for (practical) reason that exists independently of the integrated skills and habits that make up the agent. Anscombe says of practical knowledge that its exercise is "nothing but the doing or supervising of the operation of which a man has practical knowledge" and that "'[i]ntentional action' always presupposes what might be called 'knowing one's way about' the matters described in the description under which an action can be called intentional, and this knowledge is exercised in the action and is practical knowledge" (1957/2000: §48). Now, on the powers framework, successful practical reasoning backed by practical want, i.e., successful cogent practical cognizing, just is the exercise of practical knowledge. Following Anscombe, then, practical reason is nothing but the ability to order (or supervise) the operations whose execution by an agent would be an exercise of her practical knowledge.

Put another way, practical reason is just the agent's ability to order intentional action concepts. And since these concepts are domain-specific and since the agent's ability to order and execute them in accordance with an appropriate ordering is a determinate of the shape *in concreto* of some skill of hers, her practical reason is a part of every skill she has acquired. Practical reason, then, does not have an independent existence as an isolable capacity that sits atop her other skills. If it did then the structure practical reason brings to the exercise of skill would likewise exist independently of skill or its ordering intentional action concepts would be redundant or epiphenomenal. To avoid redundancy or making practical reason into an epiphenomenon, the powers framework would have to commit to an additive theory of rationality on which skills are not themselves internally calculatively structured rational dispositions for action but are capacities for action that are structured externally by practical reasoning. But then it would seem that for practical reason to do its ordering work, such ordering would have to consist in the tokening of propositional attitudes that (non-deviantly) cause and guide the agent's behavior in accordance with the order represented by that attitude. To do its work, then, practical reasoning would have to consist in the sort of thing that causalism suggests it consists in. In which case, the powers framework would again fail to be a legitimate alternative to its causalist rivals. If this is all right, then, to resist collapsing into causalism, the powers framework ought to hold fast to a transformative theory of rationality on which (at least) practical reason has no existence over and above the skills that its exercise is part of manifesting.

All that said, it is also clear that practical reason has general application to every domain of skill. That is, in human beings, for every skill, one and the same practical reason is at work in deliberating about or in exercise of a skill for doing *B* and in deliberating about or in the exercise of a skill for doing *A*. Integration and social praxes help make sense of this unity while nonetheless holding on to the idea that practical reason is nothing over and above the skills whose intentional action concepts that reason functions to order. Integration helps make sense of the unity because it consists in the amplification of an order that exists already within the skills the agent already possesses to the acquisition or mastery of novel skills. Despite being a distinct structural relation between skills, it does not add structure to the intentional action concepts constitutive of those skills. It draws on already existing structure and extends it. Social praxes help make sense of the unity of practical reason because it is part of their role in the human lifeform to codify procedural knowledge-how and, in so doing, scaffold the development of individual agents' practical reason. That is, it is a function of social praxes that they set out procedures, ways of doing things, that convey calculative structure among the doables in their domains. Granted, no set of skills as instantiable by an individual is fully integrated. Social praxes as they currently exist are much too insulated from each other to allow full integration and, with it, full unity of rational agency. But there is enough overlap among everyday skills and among sufficiently related social praxes that practical reason is readily identifiable as both part of every skill and at the same time a kind of glue holding integrated skills together.

1.4.2 Basic Action

Turning to basic action, following Small (2019), I hold that basic actions are immediately instantiated practicable means. More specifically, they are exercises of incorporating habits. The debate around basic action arguably begins with Danto (1965) and Davidson (1971) but is brought into focus by Hornsby (1980) and has been a hotly debated topic ever since. The issue surrounding basic action stems from the need to avoid a vicious regress. Intentional actions have a means-end structure. That is, to be performing an intentional action one must be doing *D* by doing *C* for some *D* and some *C*. Yet, if one aims to do *C* with a view to doing *D* such that one's doing *C* is also intentional then one must find some *B* such that one can do *C* by doing *B*. Yet, if one aims to do *B* with a view to doing *C* with a view to doing *D* such that one's doing *B* is also intentional then one must find some *A* such that one can do *B* by doing *A*. Yet, if one aims to do *A* with a view to doing *B* with a view to doing *C* with a view to doing *D* such that one's doing *A* is also intentional then… and so on. If all intentional action must have a means-end structure to count as intentional then such actions can never begin, for one

would always need to supply further means to give her current means that structure. To stop this regress, philosophers have proposed the idea of a basic action. Basic actions are thought to be simple in that they do not themselves contain a means-end structure. Returning to our master pianist, her positioning her hands thus-and-so is basic because there are no further means that she must take in order to do it and can do this independently of any given end. She can do it as a means to initiating play or in instructing another or just because she feels like it. Her positing her hands thus-and-so is a purely mechanical doing.

A number of characterizations of basic action have been offered. Some have identified basic action with bodily movements like bending one's knee or elbow, raising one's arm, walking, and so on (Setiya 2007, Steward 2009). Call this the Bodily Movement Account. Others have identified it with bodily movement or skills that are mastered such that their internal calculative structure dissipates (Hornsby 2013, Frost 2016, Wolfson 2016). Call this the Basic Skills Account. On this account, anyone who has mastered, say, the art of tying their shoes will be able to tie their shoes just like that and may struggle to parse the movements of their hands and fingers that facilitate looping and knotting the laces. They no longer know how to tie their shoes except by tying them. Yet others have identified basic action with "simple skills" that have no internal calculative structure but are specified into practicable means in relation to the particular circumstances of performance (Lavin 2013). Call this the Simple Skills Account.

Small (2019) claims that there are insurmountable problems with each of these three approaches because they all deny that basic actions have an internal means-end structure. Drawing from Thompson (2008: 107-108), Small argues that the Bodily Movement and Basic Skills Accounts fail in cases where an agent is doing *A* in order to do *B* and where doing *A* is a proper part of doing *B*. This is the so-called "initial segment" argument. Thompson uses the example of pushing a stone from one place, α, to another, β, in order to push it from α to ω. Thompson argues that in such a case, the agent is pushing the stone from α to β in order to push it from α to ω and that her pushing the stone from α to β is a proper part of her pushing it from α to ω. Moreover, in such a case, because wherever we stipulate β to be is arbitrary, at every point along the path from α to ω, we can say that the agent is pushing the stone from α to that point in order to push it to ω and that her pushing it from α to that point is a proper part of her pushing it from α to ω. Importantly, there is no initial segment or point along the path where the agent is merely intentionally pushing the stone from α. As soon as the agent begins, she is at every point pushing the stone from α to that point in order to push it to ω. Both Thompson and Small claim that the spatiotemporal density of her action is teleological in nature and, thus, that every part of the act of pushing the stone from α to ω is itself an intentional action that is internally structured by a means-end

order and performed for the sake of the whole. That is, the spatiotemporal density is just the *a priori* teleological unity of the action from its moment of initiation and ongoing performance through to its completion as a unified developing process. Thompson and Small conclude from this that every intentional action has proper parts, namely, a means and an end. If this is right, then the Bodily Movement and Basic Skills Accounts are wrong from the jump: there are no movements or exercises of skill that can count as being done intentionally in the absence of an internal calculative structure. Yet, Small argues convincingly that the initial segment argument does not address the Simple Skills Account (2019: 8*ff.*) and does not provide a satisfying alternative for stopping the regress that introducing basic action was meant to solve.

To provide such an alternative and make trouble for all three accounts, Small points to the fundamental idea of action as a constraint on philosophical theorizing about it. According to that idea, actions are changes that originate or have their source in the agent and that the agent relates to not as an observer or spectator but as an agent (2019: 8*ff.*). On the powers framework, the agent's cogent practical cognizing is both what makes her the originator or source of her intentional action and what relates her to that action as agent by providing its *a priori* teleological unity in the means-end order. The problem for the three aforementioned accounts of basic action is that they fail to capture the way that the agent relates to her actions. By claiming that basic actions are simple, the Bodily Movement and Basic Skills Accounts expunge any teleological unity from those actions and, thus, make no room for the agent to relate to her basic actions as an agent. If they truly lacked an internal means-end order, the agent could only relate to her basic actions as an observer or spectator. Her arm raising, walking, or tying her shoes would be purely mechanical changes that she could initiate or cause but which for her would be opaque happenings to be explained by cognitive neuroscience or by the science of motor learning. They would be artifacts of her will, effects of events of "a kind of pushing in another medium" that would otherwise remain completely obscure were it not for the findings of those sciences (Anscombe 1957/2000: §11).

According to Small, the Simple Skills Account fails to capture how the agent relates to her basic actions in a different way. For such accounts, basic actions are "simple skills" where to be a simple skill is to be a generic act-type that is specified in relation to the circumstances in which it is tokened. The agent may "not know any general procedure for [walking] but that nevertheless, when [she walks] on particular occasions, [she] perform[s] these particular actions by knowingly taking particular means" (Lavin 2013: 286). That is, such an agent, on any particular occasion of walking, knows herself to be walking by means of walking from here to there or around some obstacle. Unlike the other two accounts, on the Simple Skills Account, the teleological

unity and structure of action is present in the performance of basic actions. However, that teleological unity is not internal to those actions nor *a priori*. Rather, it is contained "in the material circumstances" of their performance (287). The idea is that in performing a basic action the agent relates to her action not as "I am doing *B* by means of doing *A*" but, rather, as "I am doing *B* by means of doing *B**" where *B** specifies the circumstances in which one is doing *B*.

Small argues that the Simple Skills Account misses the mark by making the teleological unity of basic actions external and *a posteriori*. The agent lacks means-end knowledge how to walk, even if she has means-end knowledge of how to walk from here to there by virtue of the material circumstances of her walking from here to there. In which case, "the generic [intentional] action concept that this particular content specifies is, for [the agent], contentless" (Small 2019: 16). And this means that, according to the Simple Skills Account, she cannot distinguish a basic act of walking from one of raising her arm or tying her shoes except by way of pointing to the distinct material circumstances in which she manifests them. But this account of how we relate to our basic actions is wrong. Human agents can surely distinguish walking, raising an arm, tying shoes, and so on from one another without any need to appeal to the material conditions in which these basic actions are instantiated and without having to relate to them as observers or spectators. We can say that walking is moving one's legs in the way constitutive of walking, that raising an arm is moving one's arm in the way constitutive of raising it, that tying shows is tying and knotting their laces, and so on. The same is true for distinguishing different manners of walking: swaggering is a manner of moving one's legs in the way constitutive of walking that exaggerates those leg movements, whereas power walking is a manner of moving one's legs in the way constitutive of walking that optimizes those movements' efficiency and involves engaging the specific muscles involved. The *a priori* teleological unity of such acts must be internal to them. That is, even basic actions must have an internal means-end structure. But, then, they cannot be simple and the threat of the regress with which we began returns.

Small's way of threading the needle between holding onto the internal means-end structure of basic action and avoiding the regress involves identifying the "*constitutive unity of acts*" (2019: 17). Whereas non-basic actions involve ordered strings of intentional action concepts that can be isolated from one another, basic actions have means that are not isolable from the ends they achieve. For instance, the unity of an agent's act of refilling the water supply by pumping is a unity of parts that can be achieved independently of one another. The agent can refill the water supply using pales and, if his stroke is too shallow, he can pump air through the pipe (Small 2012: 157). By contrast, the unity of tying one's shoes, once that ability has been sufficiently ingrained, is not a unity of parts that can be achieved independently of one another. An agent who has tied her shoes

thousands of times and is knowingly tying her shoes right now may not be able to articulate exactly how she does it but she knows that she ties her shoes by holding the laces and moving her hands and fingers in a way indexed to tying her shoes. For Small, then, the *a priori* teleological unity of basic actions has the form "I am doing *B* by doing A_{doingB}", as in, for instance, "I am walking by moving$_{walking}$ my legs" (2019: 17).

Importantly, on the powers framework, the constitutive unity of basic actions is acquired and not innate. The novice shoe-tyer who is knowingly tying her shoes knows that she is tying her shoes by holding the laces, looping them, pulling them, looping one of the laces onto itself, tying the other around it, and pulling the latter through the hole she just formed until they tighten again. What might get lost with ingraining a non-basic action into a basic action through habituation is not internal structure but, rather, the ability to isolate the means from the end they achieve. This suggests that at no point do simple, non-structured actions enter the picture on the powers framework. Infants do not start out knowing how to move their bodies only to then learn how to do things by means of moving their bodies. Rather, as Small puts it:

> My bodily movements are constitutive elements of things I have basic knowledge how to do: walk, pick up cups, turn on light switches, hit forehands, play scales, tie shoelaces, butter toast. But it is learning how to do things (that are not specified as bodily movements) that we develop control of our bodies, and those things we learn how to do are constitutive unities of acts. The idea of a "pure" arm raising is, on this view, derivative: it is an abstract union of those constitutive elements of reaching, grasping, signaling, etc. that involve raising my arm" (2019: 19).

That is, an agent's first actions are non-basic and, in all likelihood, are joint efforts partly constituted by knowledgeable acts of guidance by her caregivers. Only when she has acquired abilities to do things involving bodily movements as means is she in a position to then identify those movements for what they are. But at no point are they movements that she can make without a view to doing anything. At no point are they simple, admitting of no internal means-end order or *a priori* teleological unity. The same is true of everyday skills like tying shoes and of supposedly simple skills like walking from here to there. On the powers framework, then, all action has an internal means-end structure and, thus, an internal *a priori* teleological unity. Basic actions are simply those whose unity consists in the ordering of intentional action concepts that, so long as they are basic, cannot be severed from one another in action (and are at the very least extremely difficult to isolate in speech or thought).

I can now say how basic actions, as the powers framework understands them, are exercises of incorporating habits. And I can now say how, according to the framework, any intentional action is the exercise of at least some skill. Consider the form of practical syllogisms discussed in Appendix D:

Table 1.1 The Form of Practical Reason
On the left-hand side, depicting the informal characterization of a calculation that doing *A* (vis-à-vis *this*) is a means to doing *D*. On the right-hand side, depicting the formal characterization of that calculation.

I can do *B* by means of doing *A*, and thus	P1. [doing *A* → doing *B*]
I can do *C* by means of doing *B*, and thus	P2. [doing *B* → doing *C*]
I can do *D* by means of doing *C*, and thus	P3. [doing *C* → doing *D*]
I can do *D* by means of (doing *C*, which I can do by means of doing *B*, which I can do by means of doing *A*, which I can do by means of) **doing *A* vis-à-vis *this*.**	P4. [doing *A* vis-à-vis *this* → doing *D*]

There are two things worth noting about this way of understanding practical syllogisms. First, the intentional action concepts ordered in the table are represented as if from the point of view of an agent deliberating about what to do rather than from the point of view of an agent acting. In the latter case, the "I can do" in each description will be replaced by "I am doing." The intentional action concepts drawn on or applied in the practical syllogism remain the same despite this shift from the mental action of deliberating to instantiating the relevant order in action. Nonetheless, such a shift marks a shift in which intentional action concepts are being drawn on. For in deliberation, discursive intentional action concepts for calculating are drawn on in the formation of a discursive representation of the practical syllogism that, in turn, apply the intentional action concepts used in their contents. Second, the syllogism presented in the table represents a non-basic, actively synthesized action. This should be unsurprising given that the example from which the *A-D* order was originally abstracted was Anscombe's pumper case (1957/2000: §§23-26). The pumper is engaged in an intentional action that goes beyond any of his incorporating habits, namely, poisoning the inhabitants of the house whose water supply he is replenishing (*D*). He is likely also engaged in skilled intentional actions of replenishing the water supply (*C*) and pumping (*B*). And he is engaged in a habitual intentional action of moving his arm up and down (*A* vis-à-vis *this*). I say "likely" because it could also be the case that replenishing the water supply (*C*) and pumping (*B*) are also habitual for the pumper.

In any case, the table does not represent a single practical syllogism instantiated in action. Such a syllogism could be represented as "I am doing *B* by means of doing *A*, which I am doing vis-à-vis *this*" or [doing *A* vis-à-vis *this* → doing *B*]. Now, finally, for the main claim of this subsection: a single practical syllogism laying out a practicable means captures one of two things. It captures a skill and a habit that it subsumes and which incorporates some skill. Or it captures an incorporating habit and the skill it incorporates. We can thus characterize skills and habits in relation to the following argument form:

P1. I can do *B* by means of doing *A*.

P2. I can do *A* vis-à-vis *this*.

C. I am doing *B* by means of doing *A* vis-à-vis *this*.

Where the conclusion is tokened by virtue of the agent's instantiating the order it lays out in an intentional action of doing *B* by doing *A* vis-à-vis *this*. The practicable means as identified in premise 2, namely, "I can do *A* vis-à-vis *this*," refers to one of the agent's incorporating habits. Her taking that means, identified as "I am doing *A* vis-à-vis *this*" in the conclusion, just is her manifesting that habit in action. Such a means is something sufficiently specific for the agent and something she knows *to do* immediately once it is ordered towards the relevant end(s) and once she has recognized it through some *de agendo* representation of the particulars of the circumstances of action constituting the corresponding domain of application. The first action achievable by her implementing a practicable means as identified in premise 1, namely, "I can do *B* by means of doing *A,*" refers to one of the agent's skills. Her achieving that action, identified as "I am doing *B* by means of doing *A* vis-à-vis *this*" in the conclusion, just is her manifesting that skill and, therewith, either the habit it subsumes or the habit that incorporates it. The means-end order represented in premise 1 and the conclusion is something the agent knows how to do. That is, it is a bit of procedural knowledge-how that she possesses and which constitutes (at least) part of the shape *in concreto* of one of her skills. Her grasp of the conclusion just is the activation of that part of her skill (or the generation of the relevant intentional action acorn). And for her grasp of premise 1 to be accurate, she must actually possess the skill whose activation is that in virtue of which she represents herself as being able to do *B* by means of doing *A* vis-à-vis *this*. Finally, the conclusion drawn from premises 1 and 2 is also at the same time an exercise of the agent's practical knowledge of doing *B* by means of doing *A* vis-à-vis *this.* That is, the drawing of that conclusion is the agent's engaging in the relevant intentional action either immediately or in concluding deliberation.

Identifying the first premise of a practical syllogism with a skill, the second with a habit, and the conclusion with the well-ordered manifestation of these dispositions in action fits well with Anscombe's description of practical knowledge in §48 of *Intention*. Moreover, the power framework's account of cogent practical cognizing together with its commitments to the factualist understanding of practical knowledge and to transformative theories of rationality jointly imply such an identification. These three parts of the framework come together here. In instantiating the practical syllogism in action or deliberation, the agent's grasp of premises 1 and 2 is a function of her cogent practical cognizing in light of her possessing the relevant skills and habits. Such grasping is the activation of the relevant part(s) of those dispositions and either constitute *de*

agendo representations of the order to be affected as such or constitute embedding *de agendo* representations of the order as actively synthesizable in deliberation. That cognizing culminates in her instantiating the conclusion of the syllogism as a well-ordered intentional action by virtue of its being the full activation of those dispositions, that is, activation sufficient to initiate performance (or gestation of the intentional action acorn sufficient to sprout). And, finally, such activation and subsequent performance is an exercise of the agent's practical knowledge that she is doing *B* by means of doing *A* vis-à-vis *this*.

To see the unity in a single practical syllogism instantiated in action, consider my claim that the first action achievable by implementing a practicable means is the exercise of one of the agent's skills. If this claim were false then it would be possible for her to perform an action that was not itself internally structured by the means-end order. In that case, the powers framework would fail to capture how the agent relates to her basic actions as an agent. In order to capture her relation to her basic actions as an agent, then, every intentional action—including basic actions—must be in principle articulable in terms of the practical syllogism. The powers framework cashes this out not in terms of relations between bodily movements and truth-conditional propositional directive attitudes but, rather, in terms of the relations of subsumption and incorporation between some skill and some corresponding habit as well as their exercise in the production of intentional action. I have extended the powers framework by elaborating these relations as well as that of determination of one skill by another, and integration of skills and habits in the acquisition of a novel skill. In Sections 1.2.2-1.2.4, I explored how the mutual interdependence of skills and habits manifests in the acquisition of novel skills and habits. Here, I am exploring how their mutual interdependence manifests in the production of *any* individual intentional action.

All that said, it does not follow from the fact that intentional actions require the manifestation of both skills and habits that no one ever acts out of habit. Rather, it means that Ryle (1946/2009: 234) was only half right: not only do skills naturally contain habits, but habits also naturally contain skills. When an agent performs any non-basic intentional action, she is exercising a habit that is subsumed under some skill (what Ryle gets right). But this skill might itself be incorporated into a further habit and some skills must be so incorporated for the possibility of basic action (what Ryle misses). When an agent performs a basic intentional action, she is exercising a skill incorporated into a corresponding habit. Because I take habits in human beings to be shot through with rationality such that they too exhibit the means-end order among their constitutive intentional action concepts, I take myself to be in agreement (at least in spirit) with Small's (2019) argument concerning basic action. And because the first premise of a practical syllogism captures the internal calculative

structure of a skill—even an incorporated skill—it will still turn out that any expression of such a syllogism in action will have an internal *a priori* teleological unity and, thus, a means-end order that involves both the exercise of a habit and the exercise of a skill. At least sometimes, the instantiation of a practicable means to some end is the exercise of a skill by way of a habit that incorporates it. Such cases, like that of walking by moving$_{walking}$ my legs, involve intelligently employed habits.[61] At bottom, all intentional actions involve the manifestation of some skill. The main point, however, is that once we see how skills and habits relate to the first and second premises of the practical syllogism, respectively, we see that their well-ordered exercise is nothing other than the exercise of the agent's corresponding practical knowledge and that there are no such exercises that are not at once exercises of some skill and some corresponding habit.

1.4.3 The Life Cycle of a Skill, Again

In Section 1.2, I elaborated the life cycle of a skill by characterizing four relations within and between skills and habits. First, I introduced the notion of incorporation of skills by habits. A habit incorporates a corresponding skill when the manifestation of the former is at the same time a manifestation of the latter. In Sections 1.2.2-1.2.4, I explained incorporation in terms of the dependence of habits on the capacities for action that they manifest. The introduction of cogent practical cognition and the powers framework's characterization of basic action as modeled on the practical syllogism further illuminates that dependence. Knowing *to do*, as characterized by the second premise of the practical syllogism, does not contain an internal means-end order on its own. To have such an order and to thereby be such that the agent can relate to her action as agent, then, the relevant habits must incorporate some skill and can be subsumed by a further skill. In the former case, the manifestation of an incorporating habit is the immediate instantiation of a means-end order that the disposition embodies internally. They are basic actions like tying one's shoes or walking, where the means-end order is instantiated in action but is represented by the activation of the corresponding intentional action concepts as

[61] I am not sure if this claim puts me at odds with Small concerning basic action. Where Small and I might disagree concerns the cause of the constitutive unity of basic actions. Where Small seems to think such unity is also germane to skills, I think that, strictly speaking, it is a mark of habit. This is consistent with allowing that experts might only be able to describe their skills deictically, i.e., "I am doing that by doing this," such that, when asked to describe what they are doing in detail, their words fail them. After all, there is a difference between being able to draw on the relevant intentional action concepts in action and the ability to apply those concepts in describing the actions one is performing or had performed in order to explain them to someone else. The former type-individuate the skill and the latter is a linguistic or reasoning determinate of that skill. And while acquisition of the skill implies acquisition of such determinates, it does not imply that the skilled agent is thereby an expert at describing or explaining what she is doing or has done *to others*. Her achieving this feat is a further determination of the linguistic and reasoning determinates of her base skill.

"manipulating the laces$_{\text{tying shoes}}$ (vis-à-vis *this*)" or "moving$_{\textit{walking}}$ my legs (vis-à-vis *this*)" respectively. In the latter case, the manifestation of a subsumed habit is the instantiation of a means that forms a constitutive unity which is then related to the end provided by the subsuming skill. These are non-basic actions like tying one's shoes as part of getting dressed or walking to clear one's head. The intentional action concepts activated in such actions represent the relevant orders as "tying my shoes (vis-à-vis *this*) → getting dressed" or "walking (vis-à-vis *this*) → clearing my head" respectively. Finally, since capacities for action in human beings just are skills, this means that habits, to the extent that they exist in human beings, must incorporate some skill(s). Indeed, the need to relate to one's actions as an agent and the need for a regress stopper in the form of action that the agent need not deliberate about prior to or in the course of acting together mean that human agents must have at least some incorporating habit(s).

Second, I introduced the notion of subsumption. A skill subsumes a habit when the manifestation of the former is at the same time a manifestation of the latter. Habits that are subsumed by a skill become part of the latter's shape *in concreto* and the skills that the former incorporate become determinates of that shape. In Sections 1.2.2-1.2.4, I explained subsumption of habits in terms of the dependence of acquiring a novel skill on the habituation of would-be determinates of that skill sufficient for the agent to begin performing the action with which the skill is identified. For instance, for a novice pianist to even begin playing in a way that minimally satisfies the standards of success for the social praxis and in a way that can be shaped towards proficiency, she must be able to position herself, position her hands, coordinate them, mover her fingers independently of each other, keep time, and so on appropriately and without having to figure out how to do so on each occasion. We can now say that the manifestation of these subsumed habits involves activation of intentional action concepts that represent the relevant orders as, for instance, "sitting thus-and-so$_{\text{playing}}$ (vis-à-vis *this*)," "placing my hands thus-and-so$_{\text{playing}}$ (vis-à-vis *this*)," "moving my hands$_{\text{playing}}$ (vis-à-vis *this*)," "moving my fingers$_{\text{playing}}$ (vis-à-vis *this*)," "counting$_{\text{playing}}$ (vis-à-vis *this*)," and so on. Once subsumption has occurred, the novice might well have a *de agendo* representation of the bench as sit-downable-at-to-play, unifying these habits together under the intentional action concepts "I/one can sit down (vis-à-vis *this*) → playing."

Without subsumption, the shape of intelligent agency would look quite different. In particular, while new tasks could, in principle, be painstakingly learned with a great deal of habituation, the development and mastery of skills as well as the generalizability of structure from one skill in ampliative application to another would be severely hampered (for instance, see Agostino et al. 1996, Jessop et al. 2006, Muslimovic et al. 2007, Pendt et al. 2011). Moreover, we would more quickly lose whatever skills we had acquired. So, while, unlike

incorporation, it is not a matter of adequately accounting for basic action, subsumption nonetheless plays an essential role in accounting for the life cycles of skills and, thus, intelligent agency as we actually see those cycles play out in healthy adults.

Third, I introduced the notion of determination of skills by other skills. A skill determines another skill when the agent possesses the former and the manifestation of that disposition is a determinate of the manifestation of the latter. In Sections 1.2.2-1.2.4, I explained determination in terms of increasingly fine-grained and varied shaping of an agent's skill. The intentionality of intentional action concepts and cogent practical cognizing further illuminates such shaping. For such shaping is a matter of acquiring increasingly determinate intentional action concepts and structuring them into embodied means-end orders that the agent can then instantiate in action through manifestation of her determinable skill. And because a skill is itself a complex of such intentional action concepts so ordered, acquisition brings with it determinate shape *in concreto*.

Every skill has determinates. And every intentional action is the expression of some skill. This means that such shaping need not be as painstaking or long-lasting as my focus on cases of expert skill has made it seem. Recall our grocery shopper. Even if she does not tend to do her grocery shopping after work, her deciding to do so today makes her driving there after work a determinate of both her skill for securing food and her skill for driving. That is, in deciding thus, she activates a part of her skill for securing food that overlaps with a part of her skill for driving. If this is the first time she has grocery shopped after work, then her driving there after work involves a novel ordering of intentional action concepts constituting part of the latter skill and brought under the former skill with a view to exercising it. In deciding to go to the grocery store after work, then, she acquires a very short-lived skill for getting to the grocery store after work by driving there. This novel and short-lived skill is a determinate part of the shape *in concreto* of both determinable skills of which it is an exemplification, bringing together features of the agent's constitution and background that bring it about in other contexts that she standardly succeeds in driving from one place to another and in securing food.

Finally, I introduced the notion of integration of skills with each other. A skill integrates a habit or a different skill just in case, in coming to acquire the former, the agent gives the habit or the second skill ampliative application. In Sections 1.2.2-1.2.4, I explained integration as a part of the rational transmission of skill. We can now say that skills integrate a habit or some other skill when intentional action concepts within the domain of the latter are drawn on in manifestation of the former. Incorporation and subsumption are key aspects of the ampliative application of skills and habits from one domain to another. Moreover, integration brings to the fore the importance of cogent practical cognizing, specifically the practicality and intelligence of

such cognizing. For integration is in the first instance a matter of how skills relate to each other through the appropriate and knowledgeable use of their overlapping constituent intentional action concepts. And, often, once a given skill has been incorporated into a corresponding habit, its ampliative application can be part of the integration of meta-skills that both take that habit as subsumed and take the corresponding incorporated skill as a determinate of their shapes *in concreto*.

Playing piano is a good example of such a meta-skill. Assuming the agent has some previous experience with typing—a plausible assumption nowadays—finger independence is arguably a relevant first-order skill that is amplified in application to her acquiring the skill for playing piano. Her cogent practical cognizing her moving her thumb, index, middle, ring, or pinky fingers thus-and-so (vis-à-vis *this*) in order to play, say, a C-Major scale is an ampliative drawing on of those intentional action concepts. Assuming she succeeds, such drawing on amounts to the exercise of her practical knowledge of playing the C-Major scale by means of moving the relevant fingers thus-and-so (vis-à-vis *this*). At acquisition, her ability to play piano integrates the relevant parts of her skill for typing in part by the overlap in their well-ordered intentional action concepts. The same is true of our grocery shopper: she is integrating, for however short a period of time, the relevant part of her skill for driving with her skill for securing food. Integration, then, allows us to see the importance of incorporation, subsumption, determination and the practicality and intelligence of cogent practical cognizing to the development of skill.

I will now elaborate the life cycle of a skill a second time by adding the accounts of the intentionality of intentional action concepts and cogent practical cognizing that culminates in practical knowledge laid out in Section 1.3. The previous elaboration took four things for granted. First, it took for granted that a novice can latch onto the intentional action concepts in the domain of the skill that she is aiming to acquire. Second, it took for granted that instruction and guidance by an expert facilitates the structuring of the novice's existing habits and skills in accordance with the norms of that domain. Third, it took for granted that an agent reaching proficiency acquires discursive skills related to the domain of her newly acquired non-discursive skill that allow her to discursively represent her action. And, finally, it took for granted that an agent achieving mastery acquires yet further discursive skills related to assessing the praxis with a view to adding to its edifice. Here, I cash out these different ways that an agent relates to her skill as agent throughout that skill's lifespan in terms of the means-end order, intentional action concepts, and her cogent practical cognizing. This is distinct from the discussion in Section 1.4.2 of the way that an agent in possession of an individual skill relates, as agent, to its exercise in the execution of some intentional action. After all, the moments constituting the *a priori* unity of the

life cycle of a skill differ from those constituting the *a priori* teleological unity of a single intentional action, since the former contains the latter. That said, the ways that an agent relates, as agent, to the acquisition, elaboration, and exercise of a skill are grounded ultimately in her cogent practical cognizing.

To start, I assume that the novice possesses at least some intentional action concepts in some domain other than the target domain. I likewise assume that these concepts are sufficiently close to or minimally overlap with the intentional action concepts in the domain of the target skill. And I assume that the novice has acquired language. These three assumptions are required to account for the novice's being in a position to latch onto the latter intentional action concepts through instruction and guided practice such that she can ampliatively draw on those that she already has. Giving accounts of how a human being first acquires intentional action concepts, the ability to practically reason, and the discursive determinates of everyday skills amounts to giving an account of rational transformation. And that is a dissertation, if not a lifelong project, unto itself. I take this dissertation to provide some of the groundwork for that project. But I do not take myself to be in a position to explain such transformation itself. Rather, I want to explain the shape of intelligent agency once an agent is already in the space of reasons, as it were. Finally, I will return to using "rote routines," "intelligent tendencies," "mere competences," and "expertise" to signal the progression from unintelligent habits and skills to intelligent habits and skills.

The life cycle of a skill starts with the existence of a corresponding social praxis. That praxis consists of its practitioners as well as codified procedural knowledge-how. The intentional action concepts specific to its domain are ordered in that codified procedural knowledge-how as discursive representations of type-(i) facts. They are likewise ordered within the practitioners as the determinate sub-skills and subsumed habits composing the overall shape *in concreto* of their skill in the form of ordered intentional action concepts representing opportunities for instantiating type-(ii) facts. The novice begins outside of the relevant social praxis. She does not have the corresponding skill, although she does have some of the skills and habits that will come to form the initial shape *in concreto* of her skill at acquisition. An expert relays type-(i) facts to the novice through the use of discursive representations in instruction or through embodying corresponding type-(ii) facts in demonstration. The expert likewise relays the latter facts to the novice through guidance in practice. In doing so, the expert is communicating, providing examples of, and eliciting in the novice the appropriate means-end orders of actions within the domain of skill and as determined by the relevant social praxis. He thereby communicates and provides examples of the intentional action concepts within that domain and structures the intentional action

concepts she already possesses through guiding their activation and instantiation in accordance with the means-end orders represented as appropriate within the domain.

For instance, in instructing and demonstrating to the novice how to play scales, the expert pianist conveys strings of intentional action concepts like "One can play a C-Major scale by playing C, D, E, F, G, A, B, and C notes in that order by playing C with one's right thumb, D with one's right index finger…" And in guiding the novice in practicing scales, the expert conveys those strings as type-(ii) facts for the novice to instantiate in her cogent practical cognizing, that is, "*I* can play a C-Major scale by playing C, D, E, F, G, A, B, and C notes in that order by playing C with my right thumb, D with my right index finger…vis-à-vis *this*." Strictly speaking, the novice might not be drawing on these intentional action concepts just as such. She might well be drawing on closely related ones that she already possessed prior to the start of her training. Her activation of them and drawing on them in action is nonetheless an exercise of practical knowledge of playing a C-Major scale by, say, "pressing this key with my thumb, that key with my right index finger…vis-à-vis *this.*" The expert is guiding the activation and drawing on of those intentional action concepts on the part of the novice into an order that both can recognize as the latter's playing a C-Major scale. This is true even if the novice at this stage only recognizes the order as one of playing "*that*," as in "I am playing *that* by pressing this key with my thumb, that key with my right index finger…vis-à-vis *this*."

Many of the intentional action concepts that the novice is in the process of acquiring are composed of means-end orderings of further intentional action concepts that she has acquired through guided practice of pressing keys and the like. At bottom, she will be drawing on intentional action concepts that she possesses relative to a different domain. That is, at bottom, she will be giving ampliative application to skills and habits that she already possesses, say, in virtue of having already learned to type, prior to being brought into the social praxis. The novice can come to represent the target intentional action concepts as they are represented in the corresponding codified procedural knowledge-how through instruction by the expert. The expert will present to her discursive representations (say, in speech) that she will learn to associate with the means-end orders that he is guiding her in instantiating. Whether she comes to be able to represent the relevant intentional action concepts discursively or only latches onto them by having her pre-existing intentional action concepts appropriately ordered in direct guidance does not matter at this stage. What matters in this initial stage is that her practicing instantiates those concepts in action where such instantiation is appropriately ordered by the expert in accordance with the norms of the relevant social praxis. The novice's ability to discursively represent what she is doing in line with the codified procedural knowledge-how of that praxis matters for unguided

practice. Only then must the novice be in a position to recognize the order that she is instantiating as the order that it is relative to that praxis. Only then will that order be present in the novice's action in a way that she can structure herself. In the initial stage of guided practice, the expert supplies that order in structuring what the novice can already do towards, e.g., playing a C-Major scale. As a result, the novice develops a mere competence for playing a C-Major scale thanks in large part to the exercise of the expert's practical knowledge in *supervising* the instantiation of that order.

Eventually, guided (and unguided) practice results in the incorporation of that and related mere competences by corresponding rote routines or intelligent tendencies. The means-end order that the novice instantiates through the exercise of her mere competence for playing a C-Major scale becomes sufficiently well associated with the circumstances of performance such that any time she is in those or sufficiently similar circumstances, she can immediately instantiate that order *in situ*. That is, she now has *de agendo* representations whose complete activation would result in appropriate application of a solution system for playing a C-Major scale. No further opportunities for action need be generated or exploited in applying the relevant solution system. Her acquiring the relevant incorporating habit just is her learning to, e.g., play a C-Major scale under the relevant circumstances. It then becomes a basic action for her.

Unlike acquiring the corresponding mere competence, acquiring this habit can be wrought by a mechanism of association. While acquiring the mere competence required bringing the relevant intentional action concepts under the order codified in discursive representations of the corresponding procedural knowledge-how, acquiring the habit can be a purely inductive process that does not require additional codified procedural knowledge-how or ampliative drawing on of intentional action concepts from another domain. For acquiring the mere competence already involves association of the means-end order to be instantiated with the circumstances of action under a norm of success drawn from the relevant social praxis. That is, developing the solution system for playing a C-Major scale already involves the obtaining of the relevant need- and target-facts constitutive of the domain of its application. Prior to acquiring the mere competence, the agent's cogent practical cognizing this order and the circumstances of action issues in her performance of an intentional action expressive of skills she already possessed. After acquiring the mere competence, the novice's cogent practical cognizing that order and the circumstances of action issues in exercise of that competence. What practice resulting in development of the corresponding incorporating habit does is strengthen the connection of association to the point that the novice can instantiate the appropriate means-end order in the appropriate circumstances without having to be guided by the expert or consult the relevant codified procedural knowledge-

how before or during performance. She can then practically cognize the order as an undifferentiated unit, a constitutive unity, a "doing $A_{\text{doing}B}$ vis-à-vis *this*" or, what amounts to the same, a basic action.

After the novice has acquired a sufficient number of mere competences and after she has inculcated a sufficient number of incorporating habits, she is in a position to practice and perform the action(s) defined by the social praxis and definitional of the target skill. At this stage, the expert guides the novice in the performance of those actions. The procedural knowledge-how that the agent has so far acquired is brought together into a network of means-end orders in accordance with the norms of the praxis. Developing such a network is a hard-won achievement in the case of learning to play piano. And even if the expert has never instructed the novice, she might well have at this point developed her own discursive competences for discursively representing her own actions that align fairly closely with the codified procedural knowledge-how of the praxis. Relying entirely on the expert's guidance to provide the relevant structure through *his* discursive representing would make reaching this stage incredibly difficult. Learning to discursively represent her own actions in the absence of instruction from another at the very least makes reaching this stage much less arduous, even if, in teaching herself, the novice's discursive representations differ markedly in their vehicles from the codified procedural knowledge-how of the praxis. In any case, the network of mere competences and rote routines or intelligent tendencies that is developing at this stage comes about through mutually reinforcing sub-cycles of incorporation and subsumption that I discussed in Section 1.2.4. The novice's coming to possess the intentional action concepts constituting this network is just her learning how to play and, thus, is just her acquiring the skill to play.

At this point, the bits of procedural knowledge-how that the novice has acquired in guided (and unguided) practice are transformed by the formation of the network of means-end relations into a coherent whole with multiple tracks. Such procedural knowledge-how, as embodied in the agent, forms the ground floor of the superstructure of the shape *in concreto* of her now newly acquired mere competence. She can now cogently practically cognize playing the piano. On the metaphysical side of the transformation, this solidification of the ground floor has three aspects. First, it involves subsuming the relevant incorporating habits into the skill. Second, such solidification involves making the relevant mere competences—including those incorporated by corresponding habits—into determinates of the shape *in concreto* of her now emerging skill. Third, it involves integrating whatever skills the proficient agent previously ampliatively drew the relevant overlapping intentional action concepts from with her now emerging skill. In solidifying the ground floor of the superstructure of her newfound skill, the novice has become proficient at playing piano.

All along the way to acquiring the skill, the novice is acquiring in addition to the relevant mere competences and incorporating habits recognitional competences and rote routines or intelligent tendencies with respect to the intentional action concepts applied in the performance of the relevant action. These are the discursive competences that she is developing and that I mentioned in Section 1.2.4. They consist in mere competences for constructing and using discursive representations in deliberating within the relevant domain and themselves become determinates of the shape *in concreto* of her newfound skill. Usually, such discursive competences are imparted through instruction by the expert as part of instructing the novice to instantiate the appropriate intentional action concepts themselves. That is, while the novice is learning to play, she is usually at the same time learning the idiolect of the pianist's language at the same time from a (now) "native speaker." In cases of pure self-teaching, the novice is forming her own idiolect and, thereby, reinventing the pianist's language while she is learning to play. In either case, were she not in a position to discursively represent the intentional action concepts at work in performance at all, she would not be in a position to grasp the means-end order or the circumstances in which to instantiate that order *as such*. That is, in such a case, she would not be able to cogently practically cognize playing the piano.

This does not mean that, to grasp that means-end order, the novice must actually discursively represent it in speech, writing, thought, etc. any time she sits down to play. Rather, it means that proper grasp of the order and an ability to discursively represent that order are interdependent and come together. On the one hand, they are interdependent because the skill being formed is itself internally structured by a network of properly ordered intentional action concepts. On the other hand, for those concepts to represent the order in a way that is in principle publicly available, they must be structured by at minimum the activation of corresponding discursive skills on the part of the expert or the novice herself. Recall that the relevant order is represented in activation of the relevant intentional action concepts (or generation of the relevant action acorns) as one that any agent appropriately situated could instantiate. That is, such activation represents the order as one that is actually there whenever actions of the relevant type are being performed with the corresponding intentions. For activation of the relevant intentional action concepts to represent that order *as such*, they must be structured such that they constitute a disposition for a recognizable way of doing what they are dispositions for doing. And such structuring only comes by way of discursively representing that way of doing that thing. Yet, for such discursive representations to have their content, they must apply the relevant intentional action concepts so ordered. Indeed, assuming the now proficient agent has been instructed by an expert and has not yet discursively represented the way that she plays, the relevant order is instantiated but not articulated in her cogent practical

cognizing playing. Yet, the nature of *what* she is learning in learning how to play, namely, the intentional action concepts of the domain appropriately ordered, partly explains her concomitant acquisition of abilities to discursively represent those intentional action concepts in deliberation, expression of intention, and action description, explanation and assessment.

At the stage of proficiency, the agent can turn her practical cognizing onto incorporating her newfound skill into a corresponding habit and onto elaborating her skill by modifying or adding to its shape *in concreto*. She engages in (un)guided practice and performance of what are for her novel means-end orders and, thus, draws on what are for her novel intentional action concepts consisting in novel-means end complexes within the domain of her newly acquired skill. And since she has likely acquired her idiolect of the pianist's language in acquiring the skill to play, she is now in a position to discursively represent intentional acts of playing as well as the network of means-end relations that form the ground floor of the shape *in concreto* of her skill. At first, she will put her skill and its discursive determinates to work in (un)guided practice and performance to master the domain and, thereby, her idiolect of the pianist's language as it already exists in the corresponding social praxis. In doing so, she is shaping her skill herself, adding to the ground floor of its shape *in concreto*.

However, at some point, as she reaches expertise, she will run up against some of the limits of the extant social praxis. She will have learned everything that her teachers had to offer. The shape *in concreto* of her developing expertise will constitute a superstructure of networks of intentional action concepts ordered into means and ends that she can readily deploy in action. And yet, new means-end orders and, thus, new intentional action concepts representing those orders will be available for her. Or she will find that the relevant social praxis disagrees with some of what she was taught or with what she taught herself. Some ordering of means to end is inappropriate so far as the praxis is concerned or some intentional action concept she possesses thanks to guidance or self-teaching is not really a proper part of the domain (again, so far as the praxis is concerned). She might find this out through suffering *prima facie* failures of practical cognition that, really, doing *A* is not a way of doing *B* or through being told. In any case, she is now in a position to develop whole networks of intentional action concepts constitutive of distinct styles. She can break the rules she had previously learned and form new ones with a view to mastery. That is, she can learn new ways of doing things. Her doing so, successfully and knowingly reshaping her own skill and making it even more determinate by the acquisition of new styles, is a mark of mastery. She is adding to the edifice of the social praxis and, assuming she has a modicum of pedagogical skill, could well teach someone else what she knows how to do. She might directly guide them in

instantiating means-end orders for which she has the relevant intentional action concepts or she might instruct them through constructing and using discursive representations of those orders, applying those concepts in the process.

All along the way from proficiency to mastery, the subsumed habits and determinate skills composing the evolving shape *in concreto* of the agent's skill are becoming more intelligent. Their gain in intelligence is a function of mutually reinforcing sub-cycles of incorporation, subsumption, determination, and integration as mediated by the structuring, modification, development, and (ultimately) activation of networks of corresponding intentional action concepts. That is, these sub-cycles of incorporation, subsumption, determination, and integration implicate the agent's cogent practical cognizing. On the one hand, then, the increase of intelligence of habits is ultimately a function of the agent's ability to at once practically cognize the practicable means she takes as a basic action and her ability to practically cognize that means as itself consisting in some means-end order. On the other hand, the increase in intelligence of skills is ultimately a function of the agent's ability to practically cognize the means-end order that type-individuates that disposition with increasing cogency. The degree of cogency consists in how reliably the agent can activate the relevant intentional action concepts when instantiating the corresponding means-end order is actually appropriate. Such reliability in turn consists in decreases in the likelihood of *prima facie* failures of practical cognition and mistakes in performance and, with them, the amount of correction or respecification required for success. Such decreases make possible stringing together increasingly varied intentional action concepts as part of the agent's cogent practical cognizing the means-end order definitional of the skill.

Tying the intelligence of habits and skills to the cogency of the agent's ability to practically cognize within the relevant domain makes sense of a counterintuitive phenomenon. Namely, what are normally considered habits seem more intelligent the more readily an agent can articulate and isolate what she is doing from the relevant end in speech, writing, thought, etc. By contrast, what are normally considered skills seem more intelligent the more readily an agency need not articulate nor isolate what she is doing, the more readily she can act successfully without first deliberating, and the more readily she can act successfully without having to consult codified procedural knowledge-how. It seems, then, that the intelligence of habits and skills pull in opposite directions. If they really do pull in opposite directions then it seems odd to say that it is one property—intelligence—that is increasing in both of them as they go their separate ways.

The air of mystery is dissolved once we recognize the role played by cogent practical cognizing in increasing the intelligence of both habits and skills as part of the latter's life cycle. For, again, the intelligence of

both consists in whether the agent can instantiate the means-end order contained in them in a way that minimizes *prima facie* failure and mistakes in performance. The mark of intelligence, then, is exhibited by the cogency of the agent's practical cognizing in instantiating that order in action. The less correction or respecification involved in pulling the action off, the more one is warranted in ascribing intelligence to the habit or skill exercised by another. Warrant for such ascription is therefore tied to publicly available measures of intelligence in performance. That these measures are shared between habit and skill suggests that they are measures of the same thing.

That said, the difference runs deeper. For what the intelligence of intelligent tendencies and expertise consists in is maintenance of the internal means-end structure and sensitivity to their applicability to the relevant circumstances. Over-practiced habits that devolve into rote routines become less structured and agents start to become less sensitive to their applicability. At worst, the result is a disposition that can only be manifested as a basic action and that is manifested primarily in action slips: a bad habit. The more fine-grained means-end order that was there before dissipates until all that is left is a disposition to, say, tie one's shoes in the way one always has. And an agent in this position might well struggle to tie them in a different way that she used to know if commanded to do so. By contrast, mere competences do not (yet) have enough internal structure. At worst, the agent has a disposition that can only be manifested in the performance of the corresponding skillful action itself. She cannot, e.g., play piano well, in such-and-such a style, with gusto, etc. She never possessed the intentional action concepts required for these more elaborate, flexible performances. Or she was never able to reliably successfully string those intentional action concepts together in her practical cognizing. Or both.

So, there is a sense in which the intelligence of habit and the intelligence of skill really do pull in opposite directions. Yet, their intelligence is nonetheless unified in the objects of cogent practical cognizing, namely, the ordering of intentional action concepts that are embodied the agent's skills and habits. The powers framework's account of practical cognizing uniquely positions it to elucidate this unity as part of the life cycle of a skill. Moreover, if it is right that the intelligence of habits and skills consists in maintaining their internal means-end structure and remaining sensitive to their applicability then the framework is uniquely positioned to explain the importance of what is called "parsing" and "chunking" in the empirical and philosophical literatures on skill as well as the importance of continuous (un)guided practice after mastery (on parsing and chunking, see Verwey 2010, Verwey et al. 2011, Wymbs et al. 2012, Fridland 2019). Parsing consists in the isolation of (seemingly) natural joints in an action and dividing that action into parts along those joints. Chunking consists in connecting actions along (seemingly) natural joints into a complex. The place of the means-end order in the

powers framework's account of practical cognizing and in its account of practical cognizing as part of the life cycle of a skill together account for parsing and chunking. They are nothing other than the breaking down (parsing) and stringing together (chunking) of intentional action concepts in accordance with the norms of the corresponding social praxis. The seeming naturalness of the relevant joints is merely a reflection of the order of means and ends codified in the relevant praxis.[62]

On the powers framework, the importance of continuous (un)guided practice after mastery comes out as a byproduct of the associationist principles at work in the development of incorporating habits and the threat of atrophy and disintegration that such association brings with it. Highly skilled agents need to be able to deploy their skill as a practicable means to whatever they are aiming to do in exercising that skill. This means that they need to develop the corresponding incorporating habit. Yet, if they do not return to the subsumed habits and determinate sub-skills of that skill in (un)guided practice, they risk over-practicing the skill itself. Such over-practice will, again, result in the loss of internal structure over time. If they over-practice, these agents risk losing their grip on the relevant intentional action concepts just as one might forget how a particularly sophisticated bit of phraseology in another language can be broken down into its grammatical parts. In both cases, an agent can act or use the phrase, but they are now more likely to act incorrectly or use the phrase incorrectly. At least in the case of skill, the agent ought to continually practice the means of her skilled action in isolation of exercising the skill in full. Doing so forms a bulwark against the association between the exercise of the skill and the circumstances for that exercise becoming so strong that she can no longer isolate those means at all in her practical cognizing.

Moreover, if skilled agents over-practice their skill, the resulting loss of internal structure will in turn result in increased difficulty drawing on the intentional action concepts internal to the skill to ampliatively apply other domains. Likewise, such over-practicing will make it increasingly difficult to draw on intentional action concepts external to that skill to ampliatively apply to it. In this case, the relevant loss is less like forgetting the use of a word and more like forgetting an entire language. Part of mastering a second language is being able to fluidly identify overlapping significance across distinct grammatical and pragmatic forms. One's first language, say, English, is integrated with one's second language, say, Spanish, insofar as one has facility engaging in meaningful speech and thought across both. Of course, there is not total overlap: some expressions in the one language will not find corresponding expression in the other. Nonetheless, part of losing one's native or second

[62] While our physiology surely puts constraints on that order, it does not determine it.

language is the disintegration of that language with the other as a function of the dissolution of the internal structure of the language one is losing. The same lesson applies to skills more generally. Over-rote skills lose their internal structure and, therewith, the connections between that internal structure and the internal structure of other skills. Part of the unity of the agent quite literally disintegrates. To avoid this, skilled agents ought not only continually practice the relevant isolated means but should also continually apply that means to the exercise of other skills. Doing so forms a bulwark against the association between the implementation of the means and its being a means to the exercise of one skill becoming so strong that she can no longer give ampliative application to it in another domain.

The foregoing is a final elaboration on the life cycle of a skill. It describes in more detail the development of a skill from novicehood to acquisition to mastery by inclusion of the notions of intentional action concepts, the means-end order, cogent practical cognizing, and integration. These four notions together account for the rational transmission of skill starting in novicehood; sub-cycles of incorporation, subsumption, determination, and integration leading up to acquisition; the increasing intelligence of both habits and skill along the way from acquisition to mastery; the development of unique styles that add to the edifice of the social praxis leading up to mastery; and the need for continuous (un)guided practice once mastery is achieved. They likewise *partly* explain the discursive determinate sub-skills of the shape *in concreto* of agents' skills as they progress from novicehood to proficiency and, finally, to expertise. I say "partly" because I take it that these determinates are abilities for *mental* acts of deliberating and discursively representing one's skill, the shape of that skill, and the social praxis of which one is a practitioner. And I have yet to give the powers framework's account of specifically mental action. That said, this section lays out an account of action that elucidates the determinable of which mental and bodily actions are determinates.

1.5 Conclusion

This chapter presents an extended overview of the powers framework for action in general. According to the framework, intentional actions are developing processes constitutive of an *a priori* teleological unity of moments of activation (intentional action acorn generation), ongoing performance, and completion. Such actions are at the same time expressions of skills. Skills are rational dispositions that are internally calculatively structured. They consist in networks of intentional action concepts that are connected as means and ends. Such networks are formed by incorporation, subsumption, determination, and integration as part of the *a priori* teleological unity of their life cycle in moments of practicing, doing, and teaching. Skills also bear practical-

inferential relations to each other in virtue of incorporation, subsumption, determination, and integration and in virtue of the agent's being application accountable in line with the norms of some corresponding social praxis. And these dispositions have intentionality in virtue of their constitutive networks representing orders that are actually there whenever actions of the relevant type are performed with intentions.

An agent cogently practically cognizes such orders from some end to a practicable means whenever she activates the relevant network(s) of intentional action concepts, tokening, in the first instance, *de agendo* representations of targets as interactable-with by instantiation of the relevant order such that she satisfies some need. She at the same time represents that order objectively as one that any agent, situated as she is, could instantiate. That is, she represents, say, her car as drivable such that she can satisfy her need to secure food by going to the grocery store by driving there. And she thereby represents going to the grocery store as a way of securing food and driving as a way of getting to the grocery store *as such*. The agent likewise cogently practically cognizes such orders when she forms discursive representations of them in deliberation. The cogency of an agent's practical cognizing is a function of minimizing *prima facie* failures and mistakes in performance. Often, agents make use of perception and inference to minimize such failures and mistakes. Speculative knowledge thereby supports the exercise of practical knowledge. But neither perception nor inference do the work of cogent practical cognizing. That is, neither provide the agent knowledge of her calculation what to do nor of what she is doing intentionally, why she is doing it, or how she is doing it.

Finally, habits play a key role in the performance of intentional action and in the acquisition of skill. For basic actions are exercises of habits understood as incorporating volitional dispositions. And such actions are required to stop the regress generated by treating intentional actions as necessarily having a means-end structure where the means is an agent's "doing something else" in order to achieve her end. Nonetheless, incorporating habits have an internal means-end structure. That structure has a constitutive unity: the doing of the means is not (easily) isolable from the doing of the end. Such acts are still intentional by virtue of having an internal means-end structure. Moreover, the formation of habits is required for the acquisition of skills: to reach proficiency, an agent must be able to do the bare minimum without having to deliberate prior to or during performance every time she sets out to practice or perform. Habits understood as incorporating volitional dispositions, at least some of which issue in basic action, allow her to focus on doing more than the bare minimum. At the same time, the agent needs to be able to parse the constitutive unity of her habits in continued practice of the relevant means in isolation if she is going to reach mastery.

So much for the powers framework. In the following chapter, I extend the powers framework to give an account of mental action. Mental actions are teleologically unified developing processes and expressions of mental skills. Mental habits are incorporating volitional dispositions that manifest in the performance of basic mental actions. What unifies mental action, mental skill, and mental habit as *mental* is their being actions of or dispositions for the construction and use of discursive representations.

2.0 A Powers Framework for Mental Action

2.1 Introduction

Mental actions are things that we do with our minds. Extending the powers framework to mental action requires elaborating this description. I will be arguing in this chapter, then, that, according to the powers framework, mental actions are teleologically unified processes that develop in characteristic stages of *a priori* moments. Like all intentional action, such moments include activation, ongoing performance, and satisfaction of a need. What differentiates mental actions as mental are their being processes of constructing and using discursive representations. This characterization implies that mental actions are expressions primarily of mental skills understood as rational dispositions for such construction and use. Mental habits, in turn, are volitional dispositions for such construction and use. And at least some of these habits are dispositions for basic mental actions. Finally, I will be arguing that what ties these dispositions for the construction and use of discursive representations together is their being solution systems for problems of self-unity. Practical discursive representing consisting in deliberation, expression of intention, and action description, explanation, and assessment is concerned ultimately with structuring agency towards greater integration of the cluster of skills constituting the agent as agent. Theoretical discursive representing consisting in theoretical reasoning and speculative knowledge is concerned ultimately with coming to a coherent and true "worldview." The need to successfully unify oneself as agent and to have a coherent and true worldview are interdependent. Achieving successful integration of skills requires at minimum that the agent latch onto objective features of the world relevant to appropriately structuring her agency. And achieving a coherent and true worldview comes by way of triangulating one's own discursive representations with that of others. Intersubjective agreement concerning the use of discursive representations is a communal effort, drawing on multiple agents' discursive capacities. And the exercise of these capacities consists in action, *mental* action. It is here that the true and the good come together.

Section 2.2 sets out the account of mental action, mental skill, and mental habit according to the powers framework. The former consists in the construction and use of discursive representations. The latter are rational and volitional dispositions for such construction and use. What unifies mental skill and mental habit as *mental* is their being solution systems for the construction and use of discursive representations. What makes mental action *mental* then is its being the activation or application of such solution systems. And these solution systems are themselves unified as discursive by having the function to solve problems born of needs to orient

oneself to the world. Section 2.3 expands on three dimensions of the account. First, human perception is conceptual by virtue of being recognition-based and, thus, consists in the potentiation of a multitude of mental actions. Second, mental actions are type-individuated by the type of problem such actions function to solve and by the corresponding type of interaction with targets sufficient to solve that type of problem. Third, mental actions start out as extended, are often extended, and are only sometimes covert. Finally, Section 2.4 deals with two objections from Strawson (2003) and Wayne Wu (personal correspondence), respectively.

2.2 The Account of Mental Action, Mental Skill, and Mental Habit

2.2.1 The Account

Recall from chapter 1 that the powers framework holds the following four claims. First, intentional actions in general are primarily expressions of skills, where such skills include subsumed and incorporating habits. Second, skills and habits are rational dispositions and volitional dispositions, respectively. Third, rational and volitional dispositions consist in networks of internally calculatively structured intentional action concepts understood as solution systems that function to solve certain types of problems through certain types of interaction with targets so as to satisfy corresponding needs. And, fourth, such interaction constitutes a teleologically unified process that develops in characteristic stages comprised of *a priori* moments of activation (or intentional action acorn generation), ongoing performance, and completion.

Recall also from chapter 1 that the powers framework adopts a transformative theory of rationality according to which mindedness is a genus of which human and nonhuman animal mindedness are species. Focusing on human mindedness, such mindedness is a matter of perceptual and non-perceptual representation as well as rational action all of which involve representing in a way that implicates truth and rational inferential relations among the relevant representations. Perception in humans is, like nonhuman perception, in effect the direct presentation of targets for kinds of interaction constitutive of types of intentional action. But, drawing on Springle's framework, human perception is not *just* a matter of having the relevant *de agendo* representations for interacting with things in the perceptible environment. For human perception draws on capacities for the construction and use of discursive representations of what is perceived. For instance, a human agent perceives an apple not just as eatable or graspable but as an apple, as red, as ripe, and so on. Her perceiving the apple as both a target for action and a target for discursive representing is due to the dual activation of both her solution systems for eating and grasping it and her *discursive* solution systems for saying, writing, judging, and so on, say, that what is before her is an apple, that it is red, that it is ripe, and so on. Activation of the latter solution systems

is sufficient for conceptual perceptual content but not full-blown application of the relevant concepts. She would only apply these concepts if she were to go on to actually say, write, judge, or whatever, e.g., that what is before her is an apple, that it is red, that it is ripe, and so on. And discursive representations so constructed and used represent by providing mediated practical access to their targets through standing in or going proxy for those targets. That is, discursive representations present their objects as being of a certain kind, having certain properties, bearing certain relations to other objects, and so on in virtue of the vehicles of those representations bearing appropriate mapping relations to those objects and appropriate inferential relations to each other.

So, the powers framework holds that activation of human solution systems in perception tokens embedding *de agendo* representations of their targets that, in turn, activate but do not apply discursive solution systems for constructing and using embedded discursive representations of those targets. The activation of such embedding *de agendo* and embedded discursive representations is constitutive of perceptual representations of objects as both affording opportunities for interaction and objectively as being what they are, as having certain properties, as standing in certain relations with other objects, and so on. At bottom, the latter representations present their objects as affording opportunities for the application of solution systems for constructing and using discursive representations of those very objects, their properties, the relations they stand in, and so on. For instance, our imagined nature photographer from the previous chapter perceives the orchid as affording an opportunity to snap a photo. This perception is the activation of his discursive solution system for using his camera to construct the relevant discursive representational vehicle, in this case, the photo. He applies that solution system in taking out his camera, pointing it towards the orchid, checking his lighting, aperture, and so on, and ultimately taking the photo. The photographer thereby comes to have both unmediated and mediated practical access to the orchid. His perception of the orchid is both opportunity-for-action-based and discursive by activating solution systems for directly interacting with it on the one hand and, on the other, activating solution systems for saying, writing, judging, and so on things about the orchid as well as for snapping photos of the orchid. His application of this latter discursive solution system is his constructing a discursive representation of the orchid, namely, the photo.

Activation of human solution systems in the absence of perceiving a target tokens embedding *de agendo* representations of the vehicles of discursive representations or of opportunities to construct such vehicles. Such *de agendo* representations are the activation of discursive solution systems for using or constructing those vehicles to represent the absent target and, thereby, give the agent mediated practical access to it. Application of discursive solution systems, then, involves *de agendo* representations for interacting with or

generating vehicles for discursively representing some target to which, through use of those vehicles, the agent then gains mediated practical access. For instance, our imagined orchidologist from the previous chapter uses a photo to represent the objective identifying features of a certain species of orchid. Her representing those features of that species of orchid just is her applying the corresponding discursive solution system. And this application involves her having embedding *de agendo* representations of the relevant vehicles, in this case, the distribution of pigments in the photo. In this case, capitalizing on these *de agendo* representations depends on the distribution of pigments bearing appropriate mapping relations to orchids of the relevant species in light of the social praxis of nature photography. It is through capitalizing on her *de agendo* representations and, thereby, making appropriate use of the photo through trading on those mapping relations that she then gains mediated practical access to the species of orchid that she is studying. Such application, then, is constitutive of non-perceptual discursive representations of objects as being what they are, as having certain properties, as standing in certain relations, and so on as well as perceptual representations of the vehicles of those discursive representations.

In both cases of perceptual and non-perceptual discursive representation in humans, the appropriateness of the embedded discursive representations envelopes their veridicality conditions. They are appropriate just in case their vehicles *would* stand in the relevant mapping relations to the targets on the relevant occasion *were* the agent to attempt to indirectly practically access those targets through the use of those vehicles on that occasion. This is just as true for the construction and use of the phonetic vehicles of spoken English in actually saying of an apple that is present "that is an apple" as it is for the nature photographer's snapping a photo and the orchidologist's use of that photo. Moreover, embedded discursive representations stand in inferential relations to each other in virtue of their vehicles both bearing appropriate mapping relations to their objects and bearing appropriate systematic relations to the vehicles of related discursive representations. The formation of such systematic inter-vehicle relations is usually conventional but tied to veridical discursive representation by intersubjective agreement concerning their use and the nature of the targets they are used to represent. The symbol systems constitutive of a language are a good example of how convention and discursive social praxes intersect with the constraints involved in intersubjective agreement. Nature photography and other realist forms of art less so. In either case, however, acquisition of a solution system for the construction and use of discursive representations comes with acquisition of a battery of such solution systems. And acquiring these solution systems is a matter of being brought into corresponding discursive social praxes.

Finally, rational action has intentionality with objective import by virtue of being the activation of solution systems whose application instantiates an order that is actually there whenever actions of the relevant type are performed with corresponding intentions and that the agent recognizes this as such. With respect to the agent's recognizing a way of doing something objectively, such solution systems represent the orders instantiable by their application as ways that any agent appropriately situated could interact with the relevant target(s) so as to satisfy the relevant need(s). These solution systems represent such orders objectively by virtue of being shaped as part of the rational transmission of skill. Such transmission occurs at least in part through the tokening corresponding discursive representations of those means-end orders as part of guided practice and by concomitant activation of discursive capacities for deliberation, expression of intention, and action description, explanation, and assessment as part of unguided practice, performance, and teaching. The latter concomitant activation is at least partly constitutive of the agent's cogent practical cognizing what she is doing. That is, if she is engaged in an intentional action, she is thereby primed to deliberate about what she is doing, express her intention, or describe, explain, or assess what she is doing. On the powers framework, her readiness just is the activation of corresponding discursive solution systems. And such readiness allows her to identify and rectify *prima facie* failures of practical cognition or mistakes in performance and allows her to respecify her intention if need be. It is therefore an integral part of her cogently practically cognizing the means to her ends. According to the powers framework, then, intentional actions in neurally intact adult human beings under normal conditions with sufficient skill involve the dual activation of the solution system whose application is the performance of that action and of discursive solution systems for objectively representing that action.

Enough recap. The powers framework takes a broadly Rylean approach: it identifies mental actions with the activation or application of solution systems for constructing and using discursive representations *full stop*. Covert mental actions are activation of these solution systems sufficient for activating further corresponding recognitional discursive sub-solutions. Overt and extended mental actions are applications of discursive solution systems sufficient for activating further corresponding recognitional discursive sub-solutions. The framework thus identifies mental skills and mental habits with discursive solution systems *full stop*. According to the framework, two things tie activation and application of discursive solution systems together as constitutive of mental agency. On the one hand, the human species of mindedness is differentiated by its being thoroughly infused with discursive contents. On the other hand, mindedness in general is a matter of the *potentiation* of actions or, what amounts to the same thing, the *activation* of solution systems. As Boyle and Lavin (2010) put it "to be a rational creature means just this: to live by *thought*" (175; original emphasis).

Human mindedness specifically is a matter of the potentiation of acts of constructing and using discursive representations or, what amounts to the same thing, the activation of specifically discursive solution systems. So, if the powers framework is right that mental agency is at bottom a matter the potentiation of recognitional discursive sub-solutions as a function of at least the potentiation of the construction and use discursive representations then mental actions are pervasive but not coextensive with non-mental actions. I will say more about this later. For now, I want to further unpack the powers framework's account of mental action and mental agency.

Like non-mental skills and habits, the framework distinguishes mental skills from mental habits according to whether the relevant discursive solution systems can be applied in doing what is contrary to the type of solution that they function to constitute (skill) or not (habit). That is, a mental habit is a volitional disposition for constructing and using discursive representations. Its manifestation is subject to assessment against normative standards that are internal to that disposition and which are tied, in the first instance, to successful performance. For instance, an agent's attempt at immediately deriving the sum of two numbers can succeed or fail, and her summing can be better or worse, say, with respect to the efficiency of the equation she applies or how efficiently she applies it. Yet, if the agent's solving this math problem is really an expression of a mental *habit* of hers then, in manifesting that habit, she cannot be intentionally making a mistake in calculating. That said, this habit of hers incorporates her skill for solving addition problems.

Mental skills are rational dispositions for constructing and using discursive representations. Their manifestation is subject to assessment against normative standards that are internal to those dispositions and which are tied, in the first instance, to successful performance. An agent who lacks the relevant habit but possesses the skill might attempt deriving the sum of two numbers. She likewise can succeed or fail, and her attempt can be better or worse with respect to efficiency. Moreover, both this agent and the habitual adder manifest the same *rational* disposition for solving addition problems in showing someone else how *not* to sum or in intentionally getting the wrong answer, say, in order to feign ignorance. The difference is that the latter does this at the cost of successfully resisting manifesting her corresponding incorporating mental habit. When she does this, she impedes her habit so as to express her skill.

It is worth emphasizing that, according to the powers framework, mental actions are the activation or application of complex solution systems whose potentiation consists of tokening embedding *de agendo* representations and embedded discursive representations. The embedding *de agendo* representations are the potentiation of the application of component solution systems for the construction and use of discursive

representational vehicles. The embedded discursive representations are constituted by the application of component solution systems for the construction and use of those representations through the construction and use of the corresponding vehicles. Because there is no mediated practical access to targets in the absence of unmediated practical access to those targets that stand-in or go proxy for the former, there is no mental action in absence of at least potentiation of the construction and use of the relevant vehicles. Mental skills and mental habits, then, are rational and volitional dispositions, respectively, with opportunity-for-action-based and discursive components.

Similarly, like non-mental skills and habits, mental skills and habits are internally calculatively structured complexes of intentional action concepts. These intentional action concepts, so structured, represent means-end orders that the agent herself or anyone situated as she is and in possession of those concepts can instantiate. For instance, both our habitual and skillful adders possess intentional action concepts for summing that might well represent the corresponding means-end order as "I/One can sum *n* and *m* by means of taking the quantity represented by "*n*" and combining it with the quantity represented by "*m*" by means of using the equation "*n* + *m* = *r*"." In drawing on these intentional action concepts in the initial phase of applying their discursive solution system for summing, both practically cognize the means of summing *n* and *m* on this occasion. That is, both token an *de agendo* representations, say, of the discursive representational vehicles "*n*" and "*m*" that stand in for the quantities *n* and *m*, respectively, and which are present in their perceptible environment. These vehicles are presented to both agents as affording use in the construction and use of the equation "*n* + *m* = *r*." Such presentation just is the activation of the corresponding discursive solution system, drawing on those intentional action concepts and, thereby, representing the order as "I am summing *n* and *m* by means of taking the quantity that "*n*" stands in for and combining it with the quantity that "*m*" stands in for by means of using the equation "*n* + *m* = *r*" vis-a-vis *this*."

Their application of this discursive solution system is just instantiating that order, summing the two numbers by means of combining their quantities by means of use of the equation. And each agent's doing so, if successful, results in the construction of a vehicle "*r*" that stands in for the quantity *r*. Finally, in the good case, once they have applied the discursive solution system for summing, both will recognize "*r*" as standing in for *r* and will recognize that *r* is the combined quantity of *n* and *m*. Such recognition is a further embedding *de agendo* representation that presents "*r*" as usable in the application of further embedded discursive solution systems for representing *r*, including those for simply asserting "*r*!" or "*n* plus *m* equals *r*" and so on. That is, recognition of *r* as the answer consists in activation of the myriad ways in which one can go on to competently

use discursive representations of *r* in one's thinking or action. Put another way, such recognition consists in the agent's knowing how to go on using representations of *r*.

The foregoing example presents a single "movement of mind," a single mental action of summing. On the powers framework, any mental action consists in transitions "in thought" from at least the activation of some discursive solution system to the activation of further interrelated discursive solution systems that are tied together by virtue of being ways of interacting with the same discursive content. In the case under consideration, the mental action consists in a transition from the application of both agents' discursive solution system for summing to recognition of *r* as the answer. We can represent the structure of mental action as illustrated with the example diagrammatically as follows:

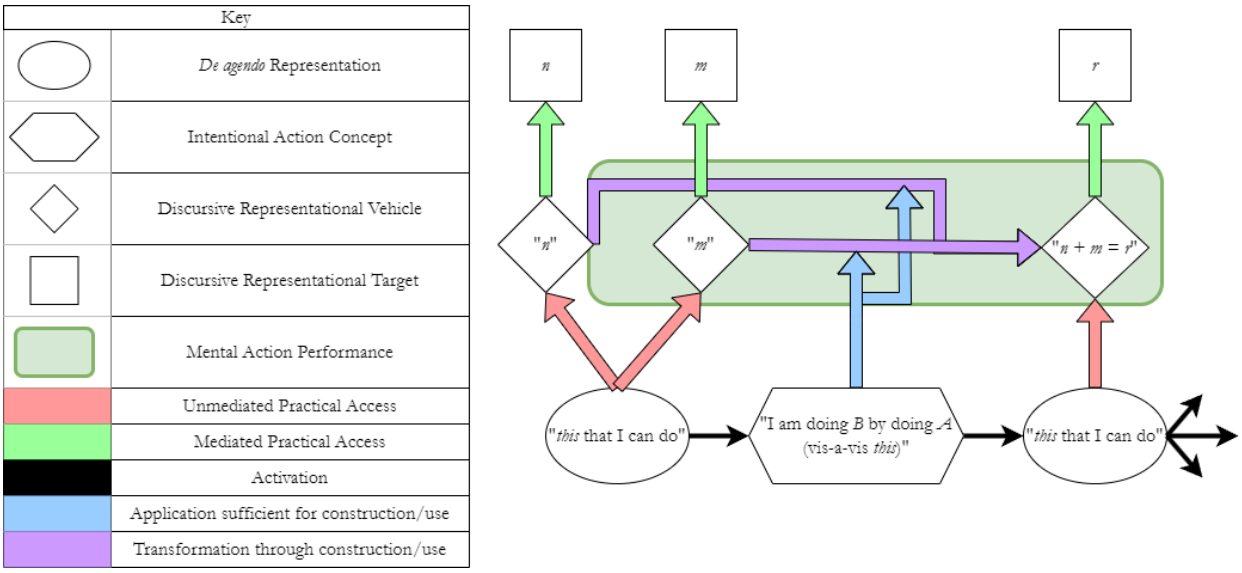

Figure 2.1 Overt Mental Action
A diagram of an overt mental action of summing two numbers.

There are at least four things worth noting about the example as depicted in Figure 2.1. First, both the skilled and habitual adders have determinate mental skills for constructing and making use of the "+" and "=" operators as well as determinate mental skills for constructing and using numerals that stand in for various quantities. For both adders, these determinate skills form part of the shape *in concreto* of their mental skill for summing, itself a determinate of the mental skill for doing arithmetic. The habitual adder differs in having incorporated these mental skills of hers into corresponding mental habits that are, in turn, subsumed under her mental skill for summing. When both adders initially acquired the mental skill for summing, they integrated the mental skills for constructing and manipulating these vehicles with their newfound skill. Second, and relatedly,

the choice of "$n + m = r$" is for convenience. It is possible that either agent constructs "$n + m$" first and only then uses this vehicle to construct "r." In that case, both "$n + m$" and "r" represent the quantity r. Or, either agent might construct vehicles for column addition. In that case, only the vehicle that they construct under the horizontal line at the bottom of the column, presumably "r", will represent the quantity r. Third, the second embedding *de agendo* representation depicted with three arrows extending from it is the activation of what I will hereafter call "a recognitional skill." Exercise of the recognitional skill in this case is the activation of further appropriate discursive solution systems for making further competent use of "$n + m = r$," "r," or both. According to the powers framework, then, such activation amounts to recognizing the quantity r as what it is. Fourth, and finally, the mental action depicted is overt. The agents are applying their discursive solution systems for summing in the actual construction and use of the relevant discursive representational vehicles so as to solve the problem of finding the sum of n and m.

But many mental actions are covert. That is, they constitute transitions "in thought" without the actual application of any discursive solution system and, thus, without the actual construction or use of any discursive representation or discursive representational vehicles. So, suppose that our agents do their summing covertly, "in their heads" as it were. We can represent the structure of such covert mental actions, extending the previous example, diagrammatically as follows:

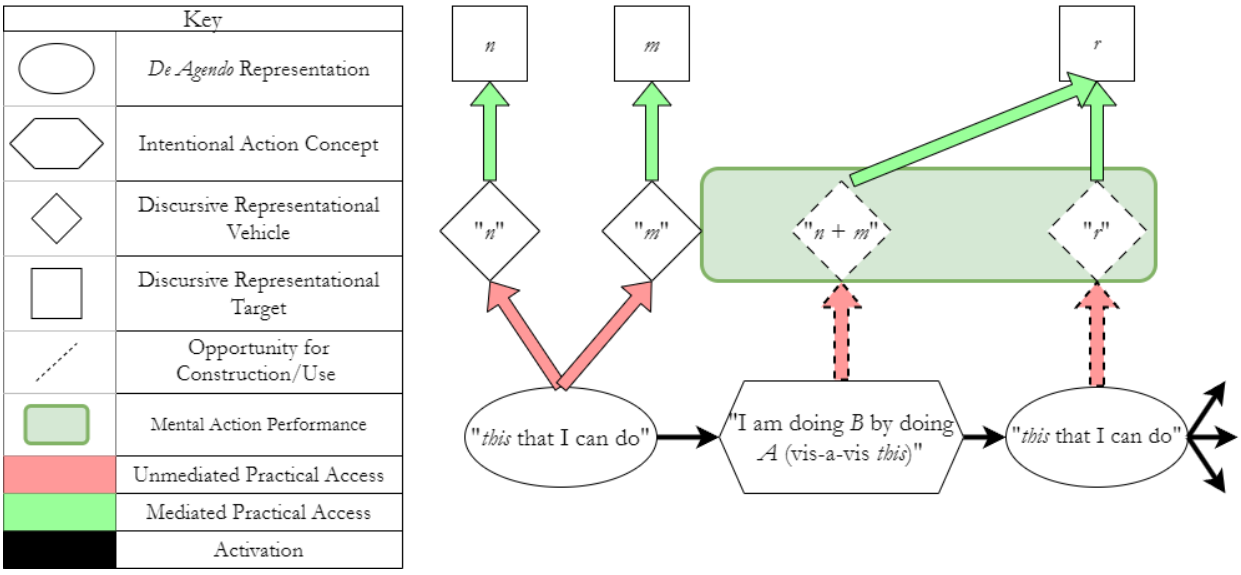

Figure 2.2 Covert Mental Action

A diagram of a covert mental action of summing two numbers.

The difference between overt and covert mental actions represented by the difference between Figures 2.1 and 2.2 consists in the fact that covert mental actions are, on the powers framework, the (mere) activation of discursive solution systems sufficient to activate corresponding recognitional skills. Figure 2.2 represents an agent covertly summing two numbers who perceives the vehicles "*n*" and "*m*" as affording an opportunity for summing and, through the activation of related networks of discursive solution systems, as standing in for the quantities *n* and *m* respectively. Her perceiving the vehicles as affording summing—her having the embedding *de agendo* representation for summing the numbers that those vehicles stand in or go proxy for—activates her mental skill summing and this, in turn, activates a recognitional skill for the quantity *r*. That is, her perceiving the vehicles activates her discursive solution system for summing the two numbers through construction and use of the relevant discursive representational vehicles. And this activation is sufficient on its own to activate a myriad of other discursive solution systems for the construction and use of discursive representational vehicles that stand in or go proxy for *r*. The activation of the agent's recognitional skill, then, is itself an embedding *de agendo* representation of opportunities for constructing and using the discursive representational vehicle "*r*" in the application of further discursive solution systems. Importantly, neither "*n* + *m*" nor "*r*" are constructed or used in this case. The agent's representing the opportunities that such construction or use *would* afford is sufficient on its own to activate the appropriate networks of related discursive solutions constitutive of the relevant recognitional skill.

Nor does an agent need to begin with perceptual representations of discursive representational vehicles in order to perform a covert mental action. She can get going simply by perceiving opportunities for engaging in covert mental actions. For instance, an agent perceiving an apple will have perceptual *de agendo* representations of the apple as eatable, graspable, and so on. Such perception, in turn, activates discursive solution systems, say, for asserting of the apple that it is an apple, that it is red, that it is ripe, and so on. And such activation can, in turn, activate non-perceptual embedding *de agendo* representations of the opportunities to assert these things as affording yet further opportunities for the application of yet further discursive solution systems for the construction and use of yet further discursive representations of the apple, apples in general, redness, ripeness, and so on. Here we have the possibility of multiple transitions in or trains of thought proper. Seeing the apple activates discursive solution systems for asserting various things of it that in turn activates recognitional skills for making further use of those assertions, say, in inferring that the apple is suitable to eat or in identifying apples in general as fruits, redness as a color, or ripeness as a state into which fruits and vegetables in general can progress. On the powers framework, such series of activation just are fully covert transitions in thought or, what amounts

to the same, fully covert mental actions. We can represent the structure of such complex mental actions illustrated by this example diagrammatically as follows:

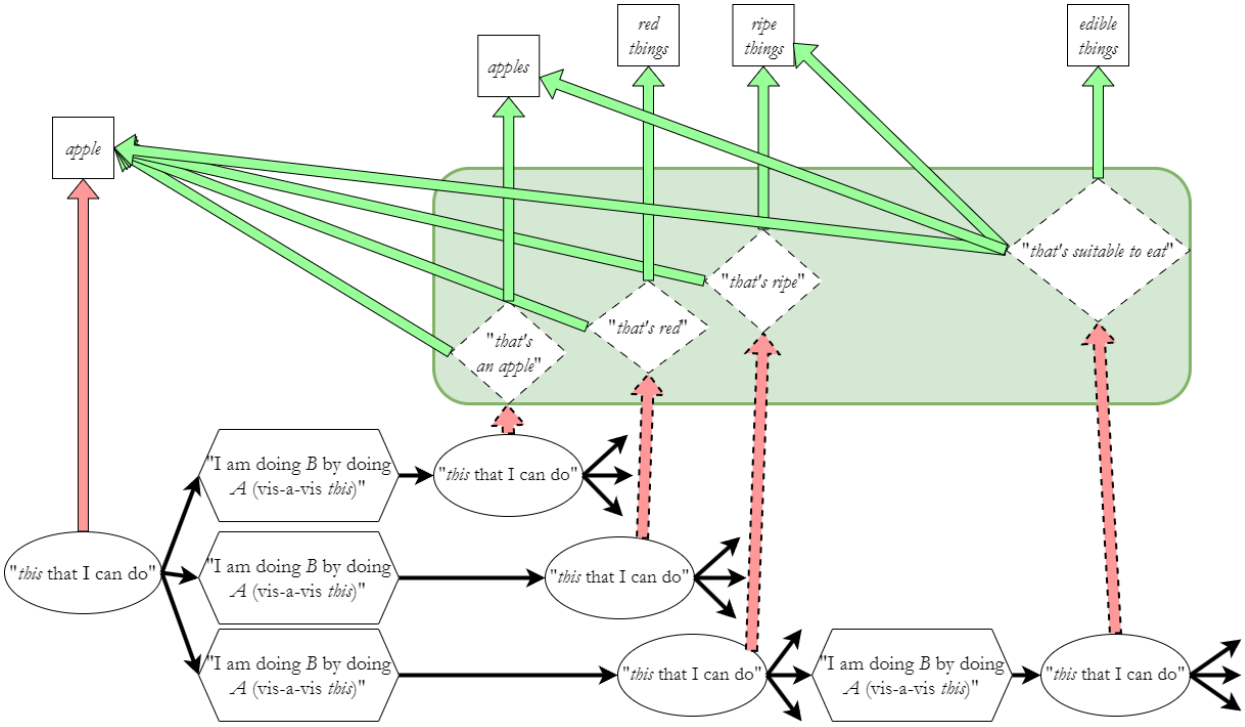

Figure 2.3 Fully Covert Mental Action

A diagram of a series of covert mental actions of recognizing of an apple that it is an apple, is red, and is ripe such that the agent then recognizes that it is suitable to eat. I am assuming familiarity with the key from the previous figure.

In the case under consideration, Figure 2.3 represents an agent's fully covert mental act of inferring that the apple before her is suitable to eat from her recognizing that it is ripe. The subsequent recognition that the apple is suitable to eat consists primarily in further activation of related discursive solution systems tied to the same content. But such recognition could be reasonably inferred from the application of related non-discursive solution systems, say, for eating the apple, for preparing it by slicing it, for giving it to someone else to eat, etc.

Recall that the powers framework holds that skills and habits in general are acquired and developed as part of the life cycle of skills. It'll come as no surprise, then, that it holds that mental skills and habits are likewise acquired and developed as part of the life cycle of specifically mental skills. There are at least three important upshots of this claim. First is that the sub-cycles of incorporation and subsumption leading to the concomitant integration and determination of newly acquired mental skills means that acquisition of such skills comes with a systematic battery of interrelated intentional action concepts for the construction and use of the relevant discursive representations. That is, in coming to possess intentional action concepts for reasoning, imagining,

remembering, perceiving constitutive of recognition, and so on, an agent thereby comes to possess a battery of systematically related intentional action concepts for the appropriate construction and use of corresponding discursive representational vehicles. Such discursive holism is unsurprising and fits well with the traditions in philosophy championed by Wittgenstein, Ryle, Anscombe, Sellars, Brandom, and McDowell, among others.

A second upshot is that such sub-cycles leading to mastery involve the incorporation of whole complexes of intentional action concepts constitutive of skills into corresponding habits that obfuscate the need for performing many relatively "easy" practicable covert mental actions by subsuming them under corresponding recognitional skills. For instance, upon seeing an apple, an adult neurally intact human in normal conditions immediately recognizes that it is an apple and, thus, a fruit. She likewise immediately recognizes that it is red and, thus, has a certain color, that it is ripe and, thus, ready to eat, and so on. In perceiving the apple, she exercises at once her mastery of perceptual and inferential dealings with middle-sized dry goods in general as well her mastery of perceptual and inferential dealings with fruit and with apples more specifically. What were once independent mental skills for inferring being a fruit from being an apple, being colored from being red, being ready to eat from being ripe, and so on during the stage of proficiency are now incorporating mental habits that are all subsumed under her skill for recognizing apples. The incorporated mental skills become determinate sub-skills of this recognitional skill. And the mental habits incorporating those determinates are activated, usually simultaneously, upon seeing an apple.

On the powers framework, such recognition potentiates the performance of a host of covert basic mental actions. That is, the agent's incorporating mental habits for discursively representing the relevant properties of the apple are already internally calculatively structured such that mere perception of the apple is sufficient to provide practicable means for attaining the relevant discursive ends. And, what is more, recalling Small (2019) from the previous chapter, the internal means-end orders for constructing and using discursive representational vehicles so as to discursively represent the apple's being an apple, being red, being ripe, and so on that are activated by the agent's merely perceiving the apple all form *constitutive unities*. Or, using Springle's metaphor, each intentional action acorn among the cluster produced by perceiving the apple contains the relevant means and ends as a constitutive unity. The means of constructing these vehicles and using them towards the end of representing the apple are not (easily) isolable from one another for this agent either in action or in (further) thought. In performing the relevant covert basic mental actions, she therefore activates intentional action concepts like "I am asserting of *this* that it is an apple by uttering$_{asserting}$ "that is an apple"," or "I am asserting of this apple that it is red by uttering$_{asserting}$ "that apple is red"," and so on. Such activation is

sufficient to further activate her recognitional skill for apples or red things by potentiating further acts of discursively representing apples in general or red things in general. Were the agent to perform these mental actions overtly, she would be drawing on these same intentional action concepts in instantiating the corresponding means-end order. That is, she would simply assert "That's an apple," or "That apple is red," where her doing so would likewise suffice to activate her recognitional skill for apples or red things by potentiating further acts of discursively representing apples in general or red things in general. Most importantly, only very little structuring of her own agency might be required for the agent to successfully instantiate the relevant means-end order in recognizing that the object before her is an apple, is red, is ripe, and so on. And, recalling the final section of the previous chapter, that only minimal structuring is at most required is indicative of mastery.[63]

And this brings me to the third upshot, namely, that practical cognizing is just as much a feature of the activation or manifestation of mental skills and mental habits as it is of the activation or manifestation of non-mental skills and non-mental habits. That is, ordering intentional action concepts from ends to practicable means in light of practical want is pervasive in the activation or manifestation of all of one's skills and habits. And this pervasiveness reveals something deep about the relationship between practical and theoretical reason. Consider that, on the powers framework, theoretical reason*ing* is a kind of discursive activity that aims at uncovering truths and which depends on the effective practical cognizing of practicable means to the ends of that discursive-*cum*-epistemic activity. Yet, as I claimed in the previous chapter in following Anscombe, practical reason can manifest in the activation of an already appropriately ordered set of intentional action concepts without the need for any actual deliberation. Practical reason is a calculation what to do from some end to some practicable means. Such calculation need not take the form of actual deliberation but can be

[63] Suppose she makes a mistake in performance and asserts that the apple is blue. In such a case, her practical cognizing is cogent but due to some mishap, she utters "that apple is blue." She need only correct her performance at this point: she might simply utter "that apple is red" or "I meant to say that that apple is red." In either case, she goes on to instantiate the order she cogently practically cognizes. Suppose she fails to know some type-(iii) fact upon which her success depends, say, that the lighting is off and the apple is actually orange. In such a case, were she to gain speculative knowledge of that fact, she would practically cognize a new order, minimally structuring her own agency towards making a veridical assertion, in this case "that apple is orange." Assertions concerning the identity and properties of middle-sized dry goods are very often the performance of overt basic mental actions. Corresponding judgments, by virtue of being activations of the very same intentional action concepts sufficient for further activation of recognitional skills, are very often the performance of covert basic mental actions. This is not to say that judgments are simply covert acts of assertion, that they are simply one's silent soliloquizing to oneself. Rather, in the case of judgment, the activation at issue is that of a multitude of ways of discursively representing something sufficient to activate corresponding recognitional skills. Judgment is being primed to assert, diagram, map, draw, perform, and so on so as to discursively represent, say, that the apple before one is orange.

embodied in the manifestation of the appropriate networks of intentional action concepts within the agent's skills and habits themselves.

So, the dependence of theoretical reason on practical reason should not be cashed out as a dependence of, say, individual acts of inference and the like on prior acts of deliberation. Some instances of theoretical reason*ing* can issue from deliberation. For instance, an agent can deliberate about which method to use in solving a math problem. Her successful deliberation can then issue in her application of the relevant method in solving the problem. But if her mental skill for solving math problems (of the relevant sort) is already appropriately internally calculatively structured prior to attempting this problem then her solving it using the relevant method will not require that she deliberate. Her cogent practical cognizing in such a case consists in the following. First, her having an embedding *de agendo* representation of the relevant discursive representational vehicles as affording use of the relevant method to solve the problem (thereby presenting the problem as such). And, second, either 1) her drawing on those concepts in applying the relevant method to find the answer *or* 2) that activation being on its own sufficient for her to recognize the answer without having to actually apply the method. Her theoretical reason*ing* in such a case will exhibit a bit of cogent practical cognizing on her part in the absence of deliberation by instantiating the appropriate means-end order either (1) in actually applying the method, say, on paper or (2) in her making the relevant transition purely "in thought."

But the dependence of theoretical reason on practical reason runs deeper. For, so far, I have only been considering instances of theoretical reason*ing*. That is, I have been focused on instances of overt or covert discursive activity aimed at uncovering truth. Recall that "the mark of practical reason*ing* is that the thing wanted is *at a distance* from the immediate action, and the immediate action is calculated as the way of getting or doing or securing the thing wanted" (Anscombe 1957/2000: §41; emphasis added). On the powers framework, the mark of practical reason is that the agent's causal powers are structured into objective means-end orders such that, were she to practically want the relevant thing and were the thing at a distance, she could get or secure or do it through the immediate action. Mental skills are themselves complexes of internally calculatively structured intentional action concepts. If this is right, it suggests that the acquisition of theoretical reason understood as mental skills specifically for epistemic discursive activities requires that practical reason already be on the scene. And while the possibility of practical knowledge of a domain depends on the relevant capacities for action being shaped through discursive representations of in principle intersubjectively available ways of doing things, the tokening of such representations need not in the first instance be the exercise of

specifically *theoretical* reason. Rather, such tokening is in the first instance the exercise of practical reason in teaching oneself and (primarily) instructing others.

To close out this subsection, recall the practical syllogism:

P1. I can do *B* by means of doing *A*.

P2. I can do *A* vis-à-vis *this*.

C. I am doing *B* by means of doing *A* vis-à-vis *this*.

Recall that such a syllogism issues in the performance of an intentional action. Recall as well that each premise corresponds either to the application of intentional action concepts in deliberation or to the drawing on of such concepts in the initial phase of an intentional action, namely, potentiation. The syllogism generalizes over intentional action. *A fortiori*, it sets out the form of the means-end order that is there whenever an intentional *mental* action is performed as well. And like non-mental actions, intentional mental actions are the exercise of at least some skill and at least some habit.

If the mental action is *non-basic* then premise 1 corresponds to the activation of a mental skill that subsumes some mental habit. Premise 2, then, corresponds to the activation of that subsumed mental habit which itself incorporates some corresponding mental skill. And the conclusion is the performance of a non-basic intentional mental action that is the exercise of the subsuming mental skill by means of the exercise of the subsumed and incorporating mental habit. If the non-basic mental action is *overt* then the exercise of these dispositions consists in their manifestation. That is, the agent manifests the subsuming mental skill by way of manifesting the corresponding subsumed and incorporating mental habit. The skilled adder in this case manifests her subsuming mental skill for summing *n* and *m* by way of manifesting her subsumed and incorporating mental habit for applying the "+" operator in applying an equation for addition. She verbally articulates, writes out, or maps out the application of the equation and from doing this asserts or writes out the result. If the non-basic mental action is *covert* then the exercise of these dispositions consist in their activation being sufficient on its own to further activate a corresponding recognitional skill. The skilled adder in this case does not manifest either her subsuming mental skill for summing *n* and *m* nor her subsumed and incorporating mental habit for applying the "+" operator in applying an equation for addition. Rather, her activation of both leads directly to an activation of a set of further dispositions for constructing and using discursive representations of the quantity *r*. In which case, she applies the equation "in her head" such that she has an embedding *de agendo* representation of the opportunities to construct and use representations of *r*.

If the mental action is *basic*, then premise 1 corresponds to the activation of an incorporating mental habit. Premise 2, then, corresponds to the activation of the mental skill which that mental habit incorporates. And the conclusion is the performance of a basic mental action that is the exercise of that incorporating mental habit. If the basic mental action is *overt* then the exercise of those dispositions, again, consists in their manifestation. That is, the agent manifests the incorporating mental habit as a constitutive unity. The habitual adder in this case manifests his incorporating mental habit for summing *n* and *m* through application of the "+" operator in applying an equation for addition where his doing the latter is not isolable from his doing the former. His verbally articulating, writing out, or mapping out the application of the equation and the result are not (easily) isolable from one another either in his action or thought. Finally, if the basic mental action is *covert* then the exercise of the agent's incorporating mental habit consists in its activation being sufficient on its own to further activate a corresponding recognitional skill. The habitual adder in this case does not manifest his incorporating mental habit for summing *n* and *m* through application of the "+" operator in applying an equation for addition. Rather, his activation of this disposition leads directly to an activation of a set of further dispositions for constructing and using discursive representations of the quantity *r*. In which case, he applies the equation "in his head" such that he has an embedding *de agendo* representation of the opportunities to construct and use representations of *r*.

The foregoing is an overview of the powers framework's account of mental action. In what follows, I will expand on three further dimensions of the account before discussing three advantages of the third dimension, dealing with two objections, and contrasting the account with standard causalist accounts of mental action. For now, I'd like to measure the powers framework up to the Neutral Account of Mental Action presented in the introduction of this dissertation.

2.2.2 The Powers Framework's Account and the Neutral Account

Recall the Neutral Account of Mental Action from the introduction:

> **Neutral Account of Mental Action:**
>
> A happening is an agent's mental action just in case:
>
> 1. that happening is the agent's forming, modifying, or maintaining some token of a type of content-bearing, personal-level state or process; and
> 2. an aim that the agent has in doing what is described in (1) is to form, modify, or maintain a token of that type of content-bearing, personal-level state or process that bears a determinate content specified by that aim.

The point of setting out the Neutral Account was to provide a shared target between the powers framework and its causalist rivals. And given how much the framework radically differs from its rivals concerning the metaphysics of intentional action and rational agency in general, the worry that these accounts end up talking past each other is live. Recall as well the standard causalist's account of mental action. According to it, a mental action is a mental state or process that is (non-deviantly) caused and guided by a truth-conditional propositional directive attitude like a desire, belief, belief-desire pair, intention, plan, goal-state, etc. Causalism about mental action satisfies the mental condition (1) of the Neutral Account by holding that what is (non-deviantly) caused and guided in cases of mental action are the formation, modification, or maintenance of content-bearing personal-level states or processes that are paradigmatically realized entirely in the activity of the agent's neurocognitive architecture. It satisfies the action condition (2) of the Neutral Account by holding that the content of the relevant directive attitude includes that the agent form, modify, or maintain the relevant state or process. So, unsurprisingly, standard forms of causalism appear to be viable as accounts of mental action as understood by the Neutral Account.

What about the powers framework? It satisfies the mental condition of the Neutral Account by holding that the states or processes formed, modified, or maintained by the agent are content-bearing and personal-level. They are content-bearing by virtue of being the potentiation or manifestation of dispositions for the construction and use of discursive representations. Such representations provide the agent mediated practical access to the relevant targets by her at least being primed to make targets to which she has unmediated practical access that stand in or go proxy for the former. When the agent tokens these representations in the performance of an *overt* mental action, she, in her capacity as an agent, forms, maintains, or modifies an actual, content-bearing discursive representation that is in principle intersubjectively available. The content of that representation has veridicality conditions and, thus, is of the sort familiar to most contemporary philosophers of mind. When the agent potentiates the construction and use of such a representation sufficient to potentiate a further, related set of discursive solution systems in *covert* mental actions, she, in her capacity as agent, forms, maintains, or modifies an embedding *de agendo* representation that embeds a corresponding content-bearing discursive representation. Yet, the latter representation is not actually constructed or used. Rather, the agent remains in the initial phase of an intentional action of construction and use. She (perceptually) represents some discursive representational vehicle, object, or the absence thereof as affording an opportunity to construct and use such a representation and, with it, an opportunity to go on appropriately using the representation she would have thereby constructed. The content of these embedding *de agendo* representations has

appropriateness conditions. Such content is realized primarily in the activity of the agent's neurocognitive architecture. Yet, such activity does not consist in the formation of content that stands in or goes proxy for anything. That is, such content is not truth-conditional. It represents by virtue of being the potentiation of whole-agent-involving interactions of discursive representing that may or may not be appropriate. Only the discursive content wrought from actual instances of such interaction is truth-conditional. That is, only when one writes, utters, pictures, etc. something where this involves actually making or using a publicly available representation does truth-conditional content come onto the scene. And this is precisely because it is only such publicly available representations that represent by standing in or going proxy for their objects. By contrast, mental representations represent by standing for actions as appropriate. Hence the content of the embedding *de agendo* representations for constructing and using discursive representations only represents by virtue of its appropriateness conditions. They represent as appropriate the construction and use of discursive representations. They do not stand in for anything. So, there are no discursive representations in the head. There are no truth-conditional contents in the head.

The content of discursive representations is of the sort that causalists claim are realized by the activity of the agent's neurocognitive architecture. This is a point of divergence between the two views: the powers framework denies that discursive representations are tokened by the activity of the agent's neurocognitive architecture. Rather, in cases of *overt* mental action, they are tokened by the agent's actual interaction with discursive representational vehicles. And in cases of *covert* mental action, the tokening of the relevant content is due to its being the potentiation of the relevant whole-agent-involving interactions. In both cases, then, the powers framework holds that the whole agent is involved in mental action. It therewith holds that mental actions are formations, modifications, or transformations of personal-level states or processes. If the mental action is *overt* then those states or processes are the in principle intersubjectively available discursive representations the agent constructs and uses. If the mental action is *covert* then those states or processes are the potentiation of the agent's dispositions for such construction and use. So, the powers framework satisfies the mental condition (1) set forth by the Neutral Account.

The powers framework satisfies the action condition (2) set forth by the Neutral Account by claiming that both overt and covert mental actions are underwritten and thereby constituted by the agent's cogent practical cognizing the relevant formations, modifications, or maintenance of the relevant content-bearing personal-level states or processes. The potentiation or application of the relevant (parts) of the agent's discursive solution systems that results in covert or overt mental actions, respectively, does not occur in the absence of the

agent's ordering the relevant means and ends. She has acquired mental skills and mental habits through the rational transmission of the former in being brought into the space of reasons. Her dispositions for constructing and using discursive representations are thus internally calculatively structured into means and ends. And when she embarks on a mental action, she is driven by practical want to identify a practicable means of achieving, securing, or getting what she wants. In cases of mental action, the agent either practically wants to form, modify, or maintain the relevant state itself or practically wants something whose achievement, securing, or acquisition can or must take as a means such formation, modification, or maintenance. In the latter case, if the relevant end is only contingently achievable, securable, or gettable by performance of a mental action, whether she performs that mental action depends on the agent's identifying it as the practicable means she will take. Identification of a practicable means in cases of mental action can take the form of a perceptual embedding *de agendo* representation of some discursive representational vehicle, some object, or of an absence of either. Or it can take the form of a non-perceptual embedding *de agendo* representation of such an absence or of an opportunity to further modify or maintain the content of a covert mental action she has just performed. And, as in cases of non-mental action, if she lapses into a *prima facie* failure of practical cognition or makes a mistake in performance and recognizes such failure, she will correct by looking for an actually practicable means or for some opportunity to correct her behavior. So long as the agent is driven by practical want to achieve something by means of forming, modifying, or maintaining some content-bearing personal-level state, then, she is aiming, through her cogent practical cognition, at forming, modifying, or maintaining that state either for its own sake or in pursuit of some further end.

The powers framework satisfies the mental and action conditions set forth by the Neutral Account of Mental Action. The framework is thus viable as an account of mental action as understood by the Neutral Account. This dissertation will have served its purpose if only this is granted. That is, I am not seeking to show that the powers framework provides the best account of mental action, one preferable to the accounts provided by its causalist rivals. The rest of the dissertation is concerned with further unpacking and defending the account provided by the powers framework. Along the way, I will register some issues I take with causalism. But I do not by any means think my scruples with causalism constitute definitive knock-down arguments. For now, I turn to unpacking three dimensions of the powers framework, which unpacking I think provides more reason to treat the framework as providing a viable alternative account of mental action.

2.3 Human Perception, Type-Individuating Mental Actions, and Extended Mental Action

There are at least three dimensions of the powers framework's account of mental action worth unpacking further. First, there is what the framework has to say about human perception. Second, there is what the framework has to say about type-individuating mental actions. And, third, there is what the framework has to say about extended mental action. I consider these dimensions in turn.

2.3.1 Perception

According to the powers framework, human perception is both opportunity-for-action-based and recognitional in nature. It is thus a kind of receptivity that draws on spontaneity. In other words, human perception is one's being "open to the way things manifestly are" (McDowell 1994: *xx*). On the powers framework, this amounts to claiming that any perceptual experience a human being undergoes is at once the activation of a host of discursive solution systems for constructing and using discursive representations of the content of that experience. And given the powers framework's account of mental action, it follows that human perception is *the potentiation of many basic mental actions*. Whether we go on to perform any of those actions *overtly* depends on whether we apply those solution systems in actually constructing or using discursive representations. Whether we go on to perform any of them *covertly* depends on whether that potentiation suffices to potentiate further such acts.

Because integration can draw multiple habits under a corresponding skill, it is an open question whether an agent's recognizing that what she is perceiving is, say, an apple, is red, is ripe, and so on is the performance of multiple basic mental actions or is merely their potentiation. To see why this is, consider a non-mental basic action of a master pianist, say, sitting down to play. Her sitting down to play forms a constitutive unity of positioning herself, her hands, her feet, and so on sufficient to begin playing. This is one action for her that nonetheless has a multitude of means that are in principle isolable from the end of preparing to play. There are multiple internal unities of means and end here that have been structured through performance and practice into a single unit that can be deployed at once should the need and opportunity arise. Unlike the novice who is just starting out, the master's intentional action of sitting down to play is not for her constituted by a multitude of further actions of assuming the right posture, assuming the right-hand positioning, assuming the right foot positioning, and so on. The same is true for the proficient pianist. And, most importantly, the same is true for proficient and masterful perceivers. Upon seeing an apple, whole networks of discursive solution systems for representing that apple are activated. The difference, then, is that the result of achieving proficiency and mastery

in perception is the possession of a skill for seeing, hearing, smelling, tasting, or touching that is often merely activated and not manifested. So, the powers framework does not hold the stronger claim that any human perceptual experience is at once the performance of many (covert) basic mental actions. It holds the slightly weaker claim that any human perceptual experience is at once the potentiation of such actions.

Marking the distinction between these claims is important. For the stronger claim implies something much less plausible than what the weaker claim implies. That is, the stronger claim implies that rational agents are all the time performing many covert basic mental actions. And this implication is implausible because it runs afoul of the distinction between passivity and activity in a way that is especially egregious. Perception, considered by itself, *seems* passive. After all, being open to the things as they manifestly are is not yet doing anything in the world. On the weaker claim, recognizing something for what it is is arguably active, but its activity is minimal. Recognition of something in perception, then, is responding to it merely by virtue of its presence. What the powers framework holds is that perception is not really *fully* or *purely* passive. So, it runs afoul of the distinction between passivity and activity but only minimally.

In a sense, nonhuman animal perception consists in recognition of opportunities for various kinds of interaction with what the animal perceives. Such recognition consists in the potentiation of actions constitutive of such interaction. Human perception, then, usually consists both in recognition of opportunities for interacting with what is perceived as well as in recognition of opportunities for discursively representing what is perceived. The latter consists in the potentiation of mental actions of discursive representation. In cases of perception in the absence of recognition, acts of exploration or discovery are potentiated. One is drawn to take a closer look or do a double take, to listen for a sound to repeat itself, to take another bite, and so on. And until one recognizes what they are seeing, hearing, tasting, and so on, they are nonetheless in a position to discursively represent that they saw, heard, or tasted *something*, though *what* it was they are not sure. So, recognition of the most general type is still present. But, in any case, the powers framework holds that perception is recognition-based, a kind of active passivity, a receptivity that draws on spontaneity. It does not hold the stronger claim that perception is pure activity or spontaneity.

2.3.2 Type-Individuating Mental Actions and the Unity of Mental Agency

The second dimension of the account worth expanding on concerns the type-individuation of mental actions. It is common to type-individuate such actions as acts of judging or inferring, deliberating, imagining, remembering, recognizing in cases of perceptual expertise, and so on. The powers framework can retain these

distinctions despite drawing their boundaries in radically different places than where a causalist might draw them. Recall from Section 1.3.1 that, on the powers framework, what type-individuates an action is primarily the type of problem that a type of interaction constitutive of the corresponding type of action functions to solve. That is, some type of need, say, the need for nourishment, makes a type of interaction, say, eating, appropriate. Interactions that function to satisfy that type of need, then, fall within the scope of ways of eating. One way to cash this out is in terms of the practical syllogism. For each type of action, there is a type of end it aims to achieve and the types of means that function to achieve that end fall within the scope of that end. Just as any argument with more than two premises can in principle be reformulated as a two-premise argument, any practical syllogism with more than two "premises" can in principle be reformulated as a "two-premise" syllogism. And just as the type of conclusion one is aiming to show the truth of makes appropriate the use of certain types of premises, the type of end one is aiming to achieve makes appropriate the use of certain types of means. The point here is that we can type-individuate mental actions by the types of end they seek to achieve and, therewith, the types of means appropriate to that end.

In chapter 3, I will argue that imagining is type-individuated by the end of completing what I there call an "imaginative project" and by the means consisting in an agent's discursively representing something appropriate to completing that project where she herself directly selects that representation's content with a view to completing the project. In chapter 4, I will argue that episodic memory ends up being a sub-type of imagining. Its end is among those subordinate to the end of imagining and its means falls within the latter's scope. Successful episodic remembering is completing the imaginative project of re-experiencing a particular event from one's personal past by means of discursively representing a scene of that event where the agent selects the content of that scene with a view to re-experiencing that event. Springle and I (under review) flesh out a full account of episodic memory according to which it is a kind of covert performative "historytelling" elsewhere. But this strategy for type-individuating mental actions extends beyond imagination. Deliberation is individuated by the end of identifying means to some other end and by the means of discursively representing an order whose instantiation by the agent in action would achieve that end. Inference is individuated by the end of achieving a true discursive representation of something and by the means of discursively representing facts about that thing that speak in favor of the truth of the discursive representation one is aiming to achieve. Perceptual recognition is individuated by the end of achieving a true discursive representation of something—usually something in the perceptible environment—as what it is and by the means of potentiating the construction and use of discursive representations of the category to which it belongs, its properties, the

relations its stands in to other things, and so on.[64] Providing full defenses of these definitions of various of our mental faculties and their activities goes well beyond the scope of this dissertation. The point for now is just to show that identifying mental action with discursive representing allows for more fine-grained analyses of various kinds of mental act.[65]

All that said, it is worth emphasizing the framework's account of the unity of mental agency. Again, what ties mental action, mental skill, and mental habit together as *mental* is the overarching function of rational and volitional dispositions for the construction and use of discursive representations. The powers framework, as I am extending it, takes a broadly virtue-theoretic approach at this point: the function of such dispositions is self-unification grounded in orienting oneself to the world. Such self-unification depends on objectively representing oneself and the world with which one interacts. The practical side of this self-unification concerns mastering skills and habits whose differential manifestation suffice to achieve specifically human flourishing. But such unity of an agent as agent is not merely a matter of coherence since the appropriateness of the application of any solution system constituting part of an agent's integrated skills inherently implies interaction with (or actual generation of) some target that would be sufficient to actually satisfy some actual need of hers. That is, for any given skill of hers, its manifestation is appropriate just in case, were she to actually have the relevant need and were there the target something that she could actually interact with (or generate), such interaction (or generation) would actually satisfy that need. Appropriateness, then, is a matter of the agent's being well positioned to interact with the world. The point of aiming at practical self-unity through the

[64] On this way of type-individuating inference and perceptual recognition, one can infer one's way to recognition. That is, one can infer, say, that apples are fruits, that Gala apples are red and both sweet and mildly tart when ripe, and so on. Such inferences impact an agent's overarching integrated network of mental skills such that, upon perceiving a Gala apple (including taking a bite), the agent is well positioned to recognize that *this* is a Gala apple. Likewise, one can recognize one's way to inference. For instance, one can perceive, say, a Gala apple and, in doing so, recognize it to be a Gala apple. One's recognition of the Gala apple as a Gala apple supports constructing and using a discursive representation of their now being in season by speaking in favor of the truth of such a representation. The agent need only activate or apply her discursive solution system for constructing and using the latter representation. If her inference is covert, what differentiates it from her act of recognition is that it occurs as a result of such recognition and not as a part of coming to recognize the thing before her as a Gala apple. But the line between recognition and inference here is blurry (as it should be).

[65] What about acts of attention? The powers framework would classify perceptual attending as perceptual recognition. In a way, this fits with Wu's account of attention as selection for action. The difference is just that selection is understood as the potentiation of a set of related discursive solution systems for discursively representing what one is attending to. Maintaining attention or paying attention, especially in contexts of interpersonal interaction, would require more than mere recognition. I am not sure what else is required. It might require consistent updating or maintaining of the agent's priority structure, whatever that amounts to. Or it might require something like continuously structuring one's discursive solution systems constitutive of (speculative) understanding. Or it might require selection-for-action of a different kind than recognition. I will leave figuring this out for later work.

construction and use of discursive representations in deliberation, expression of intention, and action description, explanation, and assessment is to shape ourselves (and others) and our actions (or theirs) towards more effective interaction (read: means) with the world and towards the correct identification of needs (read: ends).

The theoretical side of this self-unification concerns shaping as coherent and accurate a "worldview" as possible (McDowell 1994: 135). Such an achievement is fundamentally a communal effort for two reasons. First, as knowers, we are finite and depend heavily on others to fill in gaps in our epistemic take on the world. That is, we cannot hope to arrive at a thoroughgoing conception of the world that is coherent and represents the world as it is without some amount of deference to others. Second, our dependence on others outstrips such deference. For to arrive at a worldview that is coherent and veridical, we need to be able to triangulate our discursive representations with that of others if those representations are to count as objective in the first place. That is, to represent an apple not just as affording eating but as an apple, as red, as ripe, and so on it must be possible that the vehicles we are primed to construct and use bear both appropriate mapping relations to their objects and appropriate inferential relations to each other. And the possibility of competent use of such vehicles requires, in turn, that there be a set of like vehicles constituting a system, that anyone in possession of that system would be primed to use it in roughly the same way as us, and that they would use it in roughly that way were the need and opportunity to arise. Because such agreement in the use of these vehicles is not a given, some amount of triangulation in their use must occur for discursive representing to get off the ground. Such triangulation might well occur as a combined function of triangulating the targets of *de agendo* representations and the development of solution systems for mediated practical access via the establishment of certain conventions in the construction and use of sign-designs for mediating such access. In any case, to give a satisfying account of the theoretical side of self-unification, I think one needs to give an account of the bedrock of discursive social praxes constitutive of the space of reasons and, with this, a satisfying account of rational transformation. However the story of rational transformation goes, if it is one that has any hope of fitting with the powers framework then it will have to avail itself of some amount of triangulation in the appropriate construction and use of discursive representational vehicles. That is, given the constraints I have placed on the powers framework, such a story will have to appeal to bedrock discursive social praxes if it is going to avoid committing to the existence of private language or private, internal truth-conditional representations.

Finally, there are at least three points of contact between the practical and theoretical sides of continuously better orienting oneself towards the world. The first point of contact consists in the fact that

discursive representing in general is, on the powers framework, itself a type of intentional action. Again, discursive representing is the activation (intentional action acorn generation) or application of solution systems whose immediate targets are discursive representational vehicles interaction with which satisfies a need to mediately practically access some further target so as to satisfy some further need. If the agent aims to mediately practically access this further target so as to satisfy this further need then her success in accessing it depends on those vehicles actually providing such access to that further target. And those vehicles do so by standing in the appropriate mapping relations to that target and by standing in the appropriate inferential relations to each other. So, for instance, an agent might aim to mediately practically access a certain kind of berry so as to locate and eat it. But suppose the relevant berries are nowhere to be found in the agent's immediately perceptible environment and simply searching for them would be too costly. She might then settle for making use of a map of the locations of that and various other kinds of berry. For her use of the map to be appropriate, then, it must provide her mediated practical access to berries of the relevant kind by the map's depicting their location with at least a minimal degree of accuracy and by her use of it activating her solution systems for going to and searching out the relevant location for the berries that she is after. The more accurate the depiction of the location of the berries and the more competently she can use it so as to activate those latter solution systems, the more objective is her discursive representation of their location through her use of the map.

The second point of contact between the practical and theoretical sides consists in the way that speculative knowledge of type-(iii) facts buttresses cogent practical cognizing as discussed in the previous chapter. If the map our berry hunter uses is inaccurate or barely accurate or if she is incompetent or barely competent to use it, then her practical cognizing an order such as "I can find these berries by means of going to such-and-such location by means of walking there (vis-a-vis *this*)" is more liable to fall prey to a *prima facie* failure. More specifically, she is more likely to go on to attempt an action whose success depends on a type-(iii) fact of the form "these berries are in such-and-such location" that might not obtain at all or might be falsely or inaccurately represented by the map. By contrast, if the map is highly accurate and our berry hunter extremely competent in using it, then her practical cognizing such an order is much less likely to require correction or respecification in the face of *prima facie* failure. The efficacy of her structuring and exercising her agency is thus improved by her veridical discursive representation of the location of the relevant berries through use of the map. Moreover, such efficacy can increase in a graded way corresponding to the accuracy of her discursive representation of their location. The point for now is that how the agent structures her agency in light of her use

of the map improves in efficacy as a function of the accuracy of her discursive representation of the location of the berries. She makes herself a more effective agent indirectly by being a more accurate knower.

Finally, the third point of contact between the practical and theoretical side of continuously better orienting oneself to the world consists in the objectivity of intentional action concepts themselves. One the one hand, the objectivity of intentional action concepts that the agent possesses and draws on in instantiating the corresponding means-end order brings with it the obtaining of some corresponding type-(iii) fact(s) concerning the "topic of an intention" or of "exactly what is happening at a given moment (say) to the material one is working on" (Anscombe 1957/2000: §28 & §48, respectively). Such type-(iii) facts are things that the agent can have speculative knowledge of such that she really has two knowledges of the same object: the one practical and productive of the object, the other speculative and *an aid* in the production of that or further such objects. Speculative knowledge of one's actions, say, that the window is going up (because one is opening it), is of use to further cogent practical cognizing in the same way as speculative knowledge of other type-(iii) facts upon which the success of an action depends. The difference is that, in cases where the two knowledges overlap, speculative knowledge depends for its object on the successful exercise of practical knowledge.

On the other hand, theoretical activities involving discursive representations are both themselves the drawing on of intentional action concepts appropriately ordered into means and ends and can themselves be iteratively known through further such theoretical activities. For instance, asserting of an apple that it is red is an agent's activating and subsequently drawing on intentional action concepts for constructing and using a discursive representation of the apple. In drawing on these concepts in the application of the relevant discursive solution system, the agent represents what she is doing as "I am asserting that this apple is red by means of uttering "that apple is red" (vis-a-vis *this*)." And, in doing this, she primes a closely related sub-network of the very same discursive solution system for iterative expressions like "I am informing Smith that the apple is red." Her saying the latter is also her activating and subsequently drawing on intentional action concepts for constructing and using a discursive representation of her own utterance in the application of the relevant sub-network. And, in drawing on these concepts, the agent represents what she is doing as "I am asserting that I am informing Smith that the apple is red by means of uttering "I am informing Smith that the apple is red" (vis-a-vis *this*)." The appropriateness of the iterative application of discursive solution systems like this depends both on the agent's exercising practical knowledge constitutive of that application and on the possibility of the veridicality of the discursive representations that she thereby constructs and uses. Were it not possible for her discursive representations to be veridical then she would not be engaged in the epistemic act of assertion. It is

here that the proprietary of giving and asking for (justifying) reasons enters the picture. And if this is right, then it suggests that shaping our worldview is always partly a function of improving the network of intentional action concepts constitutive of mental skills and habits specifically for constructing and using *veridical* discursive representation. We've again reached a bedrock of discursive social praxes constitutive of the space of reasons.

The foregoing sketches some key features of a thoroughgoing Rylean virtue-theoretic conception of mind. On the powers framework, mental state concepts are not appropriately applied in attributing internal, private states bearing truth-conditional contents to individuals. For, according to the framework, there are no such states. Rather, such concepts are appropriately applied in attributing the activation or application of solution systems for discursive representation. They imply that an agent is competent in the construction and use of such representations, that she has been brought into the relevant discursive social praxes. The powers framework thus proposes a conception of the ground floor of the rational mind that breaks radically with much contemporary philosophy of mind. Contemporary philosophy of mind, for the most part, characterizes mental representation in human beings in terms of the occurrence of content-bearing mental events whose content bears the appropriate mapping relations to their objects and which are realized in the activity of the neurocognitive architecture. By contrast, the powers framework proposes that mental representation in human beings consists in the activation of discursive solution systems that are acquired through rational transmission as part of the life cycles of specifically mental skills. Again, an explanation of how any given agent comes to acquire their initial battery of mental skills and mental habits is an account of rational transformation. And I am not going to try to give such an account in this dissertation. I just want to acknowledge that such an account is ultimately required to do justice to the powers framework's account of mental action and, with it, its account of the ground floor of the human mind.

2.3.3 Extended Mental Action as Overt Mental Action

The third dimension of the powers framework's account of mental action worth unpacking is its claim that what makes an action mental is its being the activation (intentional action acorn generation) *or* application of discursive solution systems. "Application" here means that the agent engaged in an *overt* mental action is actually constructing and using an in principle intersubjectively available discursive representation. And this suggests that overt mental actions are by definition extended. To see why this claim constitutes a radical departure from most of the contemporary literature on mental action, recall from the introduction of this

dissertation that philosophers have claimed what makes an action mental is the absence of any corresponding bodily action (for instance, Peacocke 2009; Dorsch 2012; Wu 2023). If this were right, then doing a bit of arithmetic would only count as mental if the agent does not at any point start punching in numbers on a calculator or scribbling out the results of intermediate steps on a piece of paper. These philosophers commit to the No Bodily Movement Account discussed in the introduction of this dissertation. Others have argued that actions are mental whenever they are not in principle subject to third-person observation, that is, by being covert (Soteriou 2009; *cf.* Levy 2019). In which case, even moving one's eyes as part of visually attending to something either makes that visual attending a non-mental action or a non-mental bodily movement that is in principle separable from the mental action of attending that such movement enables (see also Vierkant 2022: 5-6). These philosophers commit to the Covert Account discussed in the introduction of this dissertation. So, the powers framework's claim that what makes an action mental is its being the activation *or* application of discursive solution systems breaks radically with standard accounts of mental action by holding that, in applying these solution systems in the performance of overt mental actions, the agent is making use of her body or material that is in principle subject to third-person observation.

Why think that application of discursive solution systems still counts as mental action beyond the claim that such solution systems in general constitute the distinguishing feature of human mindedness? Beyond recognition that, in our capacity as rational creatures, we "live by *thought*," why accept that the outward expression of thought constitutes mental action? To answer these questions, consider the following cases:

> Case 1: An agent does a bit of mental arithmetic while remaining perfectly still by imagining writing out and applying the relevant equation "in her head."
>
> Case 2: An agent does a bit of mental arithmetic by imagining writing out and applying the relevant equation "in her head" and she fidgets a bit as she is solving the problem.
>
> Case 3: An agent does a bit of mental arithmetic by imagining writing out and applying the relevant equation "in her head." She moves her finger as though she is writing the equation out and her doing so results in corresponding changes to the imagined application of the equation.
>
> Case 4: An agent does a bit of mental arithmetic by imagining writing out and applying the relevant equation "in her head." She scribbles the result of an intermediate step in her calculation so that she will not forget it when she needs to use it in solving the problem.

Case 5: An agent does a bit of mental arithmetic by imagining writing out and applying the relevant equation "in her head." She scribbles down the results of each intermediate step in her calculation so that she will not forget them when she needs to use them in making the next step.

Case 6: An agent does a bit of arithmetic by writing out and applying the relevant equation with pen and paper.

Most philosophers would allow that Case 1 is a case of mental action and would deny that Case 6 is a case of mental action. I will now try to show that Case 6 is a case of mental action and, so, continuous with Case 1. First, I contend that Case 2 does not significantly differ from Case 1. The fidgeting is bodily movement that co-occurs with the mental action. Assuming that such fidgeting is happenstance or an accidental byproduct of her mental arithmetic, both the powers framework and rival accounts of mental action would agree that such fidgeting is not part of the action. But the reasons for saying this differ. On standard accounts, the reason is that such fidgeting is bodily or is in principle subject to third-person observation. On the powers framework, the reason is that such fidgeting is not part of the means-end order the agent is instantiating in solving the problem "in her head."

To see why the latter reason is preferable, consider the claim that Case 3 does not significantly differ from Case 2 with respect to identifying what action is being performed but does significantly differ from Case 2 with respect to the relationship between the bodily movements and the mental action. There is in both Case 2 and Case 3 one action being performed: the doing of a bit of mental arithmetic. Yet, it is at least intuitive that, upon seeing the agent in Case 3 and being asked what she is doing, any normal observer in normal conditions would say "she is figuring something out." Moreover, if such a person observed more of the context of the action, he might recognize that she was doing a bit of mental arithmetic. The powers framework would explain how the observer can come to recognize what the agent is doing by claiming that her hand gestures are themselves part of the mental action by facilitating transitions "in thought" between steps in the agent's imagined application of the equation. The observer, then, is in a position to literally *observe* the agent in the midst of figuring something out, namely, the solution to the relevant math problem. And he can come to recognize that she is solving that problem in much the same way he can come to recognize that another agent is baking bread while at the moment laying on the couch by observing, say, a bowl containing dough that is covered in saran wrap.[66]

[66] With respect to the covert-overt distinction specifically, it's worth noting that cognitive scientists studying what is sometimes called "cognitive control" often describe subjects' binding an arbitrary input with both a mental representation

By contrast, consider the most common alternative: standard causalism. Recall that standard causalism claims that intentional actions are bodily movements that are (non-deviantly) caused or guided by truth-conditional propositional directive attitudes like desires, beliefs, belief-desire pairs, intentions, plans, goal-states, etc. Standard causalist accounts of mental action would have to say either that the agent in Case 3 is not performing a mental action at all because some bodily movement is involved or that the mental action she performs occurs separately from the bodily actions of making the relevant hand gestures. Insisting that her solving the equation is not a mental action in such a case because it involves some amount of bodily movement fails to mark the actual difference between Case 2 and Case 3. And claiming that there are at least two actions—the bodily action of moving her finger and the mental action of doing a bit of mental arithmetic—fails to countenance the connection between the agent's gestures and what she is doing by means of those gestures.

The causalist can attempt to recover that connection by stipulating that the gestures are "productive" rather than constitutive means to her performance of the mental action (Kelley 2024: 107*ff.*). That is, they count as means to performing the mental action by being a partial (efficient) cause of the relevant transitions in the agent's thoughts in accordance with the content of the relevant desire, belief, belief-desire pair, intention, plan, goal-state, etc. But saying this gets the wrong result: the connection between the gestures and solving the problem is not of the same kind as the connection between, say, an agent's figuring out when the Nazi leaders will arrive at such-and-such location at such-and-such a time and then setting out to pump poison into the water supply of that location shortly after their arrival. That is, the connection is not one between the obtaining of type-(iii) fact, in this case, the Nazis' arrival, which the agent then treats as such in exploiting it in the performance of a further, subsequent action whose success depends on its obtaining. Rather, the gestures are *constitutive parts* of the mental action of solving the math problem. They are constitutive means. And this is consistent with allowing that cause and effect are involved here, namely, at the level of "physical description" (Anscombe 1957/2000: §7). That is, the gestures considered purely as physical movement have among their effects corresponding changes in the activity of the agent's neurocognitive architecture. But a full description of

---

and a task-relevant behavioral response in the lab as a bit of skill acquisition (for instance, see Taatgen 2013). The skills they have in mind appear to include the task-relevant response, which can take the form of pushing a button or lever or saying a word. Yet, the skills themselves are thought to be mental. This is because, in such cases, the observable responses have no other significance in the lab other than to report the agent's tokening some mental representation rather than another. If this is right then mental actions do not have to circumvent observation, that is, they do not have to be covert to count as mental. They can include observable key-pressing and the like at least when these behaviors are meant to be (no more than) indicative of the controlled formation of task-relevant states. Case 3 is simply an extended form of the sort of thing these scientists are observing in their labs. So, those wishing to be consistent with current cognitive science should allow that Case 3 does not significantly differ from Cases 1 and 2 with respect to the action being performed.

these causes and effects would not give us a description of her action. For such a description would not characterize her gestures as means. The causalist who claims that the gestures are productive means in this case vacillates between two distinct descriptions of the same happening and hopes to hold them together. By contrast, the powers framework fixes on the description of the happening that employs "concepts of human action" (Anscombe 1957/2000: §46). And by avoiding drawing the line between mental action and non-mental action along the mental-bodily line or the covert-overt line, the powers framework gets the right result here: the gestures are part of how the agent solves the math problem.

For the same reason, the framework can claim that there is no significant difference between Case 4 and Case 3. When fixed on the description of what is happening that employs concepts of human action, writing down the result of an intermediate step turns out not to be a partial cause of the agent's moving on to the next step for the same reason that hearing the kettle whistle is not a partial cause of her taking it off of the stove. In both cases, the agents' moving on to the next step in recognition that the previous step has been completed is a function of their each having *de agendo* representations of the completion of the previous step as an opportunity to take the next step. The tea-maker relates to the kettle whistling as the obtaining of a type-(ii) fact, namely, "I am boiling water by means of heating the kettle on the stove vis-a-vis *this*," which she makes part of another type-(ii) fact, namely, "I am preparing to steep the tea by means of boiling water by means of…vis-a-vis *this*." Her cogent practical cognizing—that is, her calculation of practicable means to the end of making tea—provides the teleological unity of the ongoing process of making tea. And her responding to the kettle whistling by taking it off the stove is thereby a constitutive part of her making tea. That is, her response *is*, in these circumstances, her preparing to steep the tea and her preparing to steep the tea *is*, in these circumstances, her making tea (*cf.* Anscombe 1957/2000: §§22-23, 26).

The same is true of the agent doing mental arithmetic in Case 4. What she has written down as the result of an intermediate step affords moving on to the next step. She relates to what she has written down as the obtaining of a type-(ii) fact, say, "I am summing *n* and *m* by means of summing *r* and *s* vis-a-vis *this*," which she makes part of another type-(ii) fact, say, "I am summing *k* and *l* by means of summing *n* and *m* by means of…vis-a-vis *this*." Her cogent practical cognizing likewise provides the teleological unity of the ongoing process of mental arithmetic. And her responding to what she has written by going on to the next step is thereby a

constitutive part of her solving the relevant math problem. That is, her response *is*, in these circumstances, her going on to take the next step and her doing this *is*, in these circumstances, her solving the math problem.[67]

For the very same reason, Case 5 does not significantly differ from Case 4. In Case 5, the agent is simply making writing down the result of each step as she goes a constitutive means of solving the math problem. To insist that there is a difference between writing down one result on the way to solving a problem and writing down every intermediate result such that only the former counts as a constitutive part of a mental action is to insist on an arbitrary distinction between mental and non-mental action. Someone wishing to resist where I am going will resist the move from Case 2 to Case 3 or from Case 3 to Case 4. But once it is allowed that Case 4 counts as an instance of an act of mental arithmetic, there is no reason to think that Case 5 will not also count as such an instance. That more bodily constitutive means are taken is no reason against counting those bodily happenings as constitutive means.

Finally, Case 6 does not significantly differ from Case 5. To think otherwise is to beg the question against the powers framework and in favor of either the No Bodily Movement Account or the Covert Account. That is, unless there is independent reason to think that either the bodily-mental boundary or covert-overt boundary is significant. But there are good reasons to think that neither are significant. First, if we accept a reductive causalist approach to action that treats actions as (non-deviant) behavioral effects of certain psychological causes and treats those psychological causes as having a reductive base in the neurocognitive architecture then all actions will be effects of "brain movements" as it were. The only boundary between "internal" and "external" left on such a view is the boundary of the skin or skull. But there is no good reason to think that such a boundary draws a corresponding metaphysical boundary between mental and bodily action. At most, it draws an intuitive, conventional, pre-theoretic one. And this alone is not enough to avoid accepting that Case 6 really does not differ significantly from Case 5 once it is accepted that Case 5 is a case of mental

[67] The powers framework does not countenance productive means. After all, the arrival of the Nazis is not among the poisoner's *means* of poisoning them despite being a partial cause of her pumping the poison into the water supply at such-and-such a time. And in the previous chapter, I said that the role of type-(iii) facts in practical cognizing was only to buttress the effective calculation of practicable means to ends. The agent does not represent those facts as such in her calculation. She only represents type-(i) or type-(ii) facts in her calculation. There is no room for productive means here: either the obtaining of a type-(iii) fact is a partial cause of the agent's doing something because it supports her doing it or it's treated by her as the obtaining of a type-(ii) fact that comes to at least partially constitute her making true some further type-(ii) fact. Whether this is a weakness of the view, I am not sure. But, in its defense, I doubt that agents relate to the means that they take as causes of the ends that they achieve. Even in cases of action that extend over a long period of time or whose success is not a matter of course for the agent, use of "I am going to do *B* (by means of doing *A*)" does not, I think, reflect thinking of doing *A* as a (partial) cause of doing *B*. Rather, I think it reflects the agent's understanding that doing *B* is far off in the future or that doing *A* is not (always, usually, or in this case) a surefire *way* for her to do *B*.

action (for a similar argument, see Levy 2019: 973-975). Second, even if one wishes to avoid reductive causalism and wishes to hold onto the absence of bodily movement or covertness as the distinguishing feature of mental action, one runs into the issue of extended mind and extended mental action. There are strong arguments favoring the view that the mental can extend beyond the body. Such arguments also trade on there being no significant difference between an epistemic activity that is covert and an "epistemic action" that is the overt "alter[ation of] the world so as to aid and augment cognitive processes such as recognition and search" at least when it comes to the "spread of *epistemic credit*" (Clark and Chalmers 1998: 8; original emphasis. *Cf.* Kirsh and Malgio 1994; see also Levy 2019: 975-983).

Drilling down a bit, arguments for the extended mind tend to be similar to companions in guilt arguments. The thought is that if, say, solving a math problem or figuring out a move in a game purely "in her head" is sufficient to credit the agent with an epistemic achievement then use of the body or extra-bodily instruments to solve a math problem or figure out a move in a game is likewise sufficient to credit her with an epistemic achievement. To make this claim plausible, proponents of the extended mind set out criteria on the use of the body or extra-bodily instruments whose satisfaction are supposed to be individually necessary and jointly sufficient to count such use as the performance of an epistemic action which, if successful, constitutes an epistemic achievement. Clark (2010: 46) proposes three such criteria in the case of extended memory, which I translate into the powers framework's vocabulary for ease of exposition:

1) **Constancy:** the bodily movement or use of the instrument must be "reliably available and typically invoked;"
2) **Endorsement:** the content of the relevant discursive representation "should be deemed [at least] as trustworthy as something retrieved clearly from biological memory"; and
3) **Accessibility:** that discursive representation "should be easily accessible as and when required."

These three criteria are supposed to establish parity in the functional role of the mechanisms involved in non-extended covert mental acts of remembering and those involved in extended overt mental acts of remembering. Recalling Clark's and Chalmers's original example from the introduction of this dissertation, Otto's use of a notebook to retrieve the location of the Museum of Modern Art (MoMA) is only an instance of his remembering the location if his use of it plays the same functional role as Inga's "biological memory" when she recalls that information covertly. In which case, Otto must use his notebook any time Inga would use her

memory (**Constancy**), must treat the contents of the notebook in roughly the same way Inga would treat her mnemic content (**Endorsement**), and must be able to access the contents of the notebook as easily and reliably as Inga can access the content of her memory (**Accessibility**). If Otto uses his notebook constantly enough, endorses what he reads in it at least as often as Inga endorses her mnemic content, and can access the content of the notebook at least as easily as Inga can access her memory then Otto's use of the notebook to find the location of MoMA counts as his remembering its location.

The point generalizes beyond memory. The thought, then, is that if we acknowledge that successful covert epistemic activity culminates in an epistemic achievement then we ought to acknowledge that successful instances of the use of the body or of some extra-bodily instrument which plays the same functional role as that covert epistemic activity likewise culminates in an epistemic achievement. Any philosopher who accepts functionalism about the mind, then, should accept that the mind can be extended. The powers framework and its rivals alike accept functionalism. The powers framework accepts it in the form of Springle's problem-solving account of intentional actions. Representations, mental and otherwise, are what they are by virtue of having the function of providing practical access. That is the functional role that all representations play. Causalism accepts functionalism at least insofar as its proponents conceive of intentions (or other truth-conditional propositional directive attitudes) as being defined (at least partly) by their having the function of (non-deviantly) causing and guiding behavior constitutive of intentional action. So, both views should countenance the possibility of extended mental actions. At this point, someone wishing to resist extended mental action must fall back on either a mental-bodily or covert-overt distinction. But, as we have seen, falling back on either leads to problematic explanations of what is going on in Cases 2-6.

The powers framework has an advantage over causalism in this respect, especially once we consider the following question: why should we think epistemic creditworthiness is indicative of something specifically *mental*? The causalist has an answer that I think is wanting. Namely, in paradigmatic cases, epistemic creditworthiness is a function of the well-functioning of epistemic mental faculties. That is, in paradigmatic cases, an agent improves their epistemic standing by updating their beliefs in light of considerations of reasons or evidence. In the case of mental action, through attentive perception, inference, epistemically useful imagining, or memory, the agent gets closer to the truth. What distinguishes such actions as epistemic is their aiming at the formation of truth-conditional propositional attitudes that in fact represent the world veridically. And while such actions can be aided by things external to the agent and sometimes depend on the agent's latching onto things in the environment, these actions themselves are supposedly non-bodily or covert. In cases

of extended mental action, the agent's bodily actions or use of extra-bodily tools is (non-deviantly) caused or guided by an intention to, e.g., figure out the location of MoMA and such action or use in turn causes an update in the agent's beliefs concerning its location. Assuming such updating constitutes an improvement in the agent's epistemic standing concerning the location of MoMA, the extended mental action constitutes an epistemic achievement. That is, the action constitutes an epistemic achievement by virtue of successfully producing a propositional attitude that veridically represents some state of affairs, in this case, the location of MoMA. So, according to the causalist, if extended mental actions like this are *mental* actions, it is because the kind of happening at issue is derivative of the paradigmatic case. Extended mental acts are complex actions that take bodily or overt actions as productive means and are driven by an intention to improve one's epistemic standing through taking those means.

But why should we think that extended mental actions are mental by virtue of being derivative of the paradigmatic cases? And, more importantly, why should we think that epistemic achievements are uniquely mental? Nonhuman animals surely learn. They explore and discover things as a result of such exploration as well. And they can figure things out by "trial and error in perception" (Millikan 2006: 119). That is, according to Millikan, at least some nonhuman animals can work out what affordances there are for them in a certain situation by perceptually surveying that situation before acting and trying out different ways of solving the problem before settling on an appropriate solution. But such learning need not be understood in terms of updates to truth-conditional propositional attitudes. On the powers framework, they can be understood in terms of relatively enduring changes to their networks of solution systems in light of the opportunities for action that they pick up on or generate. And nonhuman animal exploration, discovery, and perceptual trial and error can be understood in terms of activations or applications of solution systems that result in relatively enduring changes in other such networks.

In cases of exploration and discovery, the animal might well see a space as affording exploration, go on to apply their solution system for exploration, and discover some new opportunity for action. Discovery of this new opportunity would be the activation of a yet further solution system. And the resulting relatively enduring change in the animal's network of solution systems in such a case would be the formation of a connection between the animal's solution system for navigating the space and their solution systems for making use of whatever they discovered in that space. In the case of perceptual trial and error, the animal might well see, say, the ledge of a deck as affording jumping, where the representation of such an opportunity, consisting itself in activation of the animal's solution system for jumping, in turn activates a solution system for grabbing onto the

ledge of a bird feeder some distance away. This latter activation in turn is the tokening of a *de agendo* representation of the bird feeder as grabbable-onto. The tokening this *de agendo* representation can then result in a change in the network of solution systems for jumping and grabbing: what drives activation of the former can change from some expanse of the ledge of the deck to a particular spot and what drives activation of the latter can change from the bird feeder itself to a particular location on it. In this case, the change is in what counts as a target-fact for these solution systems. And, in any case, the resulting changes seem like epistemic achievements on the part of the nonhuman animal. The animal figures out how to get to the bird feeder from the ledge of the deck. So, why is not its doing so a mental action?

One might be tempted at this point to deny that it is a mental action because such changes do not consist in the construction, maintenance, or transformation of truth-conditional propositional attitudes. But such a claim seems arbitrary at best and speciesist at worst. Considered in abstraction of our ability to construct and use in principle intersubjectively available truth-conditional representations and the discursive social praxes that enable them, why think that only humans trade in such representations? If one accepts that such acts of exploration, discovery, or perceptual trial and error are mental actions then one can hold either that animals too have truth-conditional propositional attitudes that they can agentially form, maintain, or transform or that whole hosts of actions that lead to relatively enduring changes in *human* solution systems then count as extended mental actions. Going the former route raises the question why nonhuman animals do not in general have capacities for language, map use, or other myriad discursive representational systems that seem unique to human beings. If such animals not only have truth-conditional propositional attitudes but can make agential use of them, why then do not they engage in the kind of symbol use that we do? Why do not they make external maps, express themselves through the written word, and so on? Going the latter route would have it that drilling and other forms of habituation that seem not to be cognitive achievements at all nonetheless count as extended mental actions.

Both routes seem to lead to bad results: a risk of over-intellectualizing nonhuman animals or a risk of deflating mental action and human epistemic achievement. For at least part of the point of distinguishing mental action from non-metal action is to mark a significant distinction between actions that result in epistemic or imaginative (or, more broadly, cognitive) achievements and those that do not. And, surely, a constraint on successfully marking this distinction is that one's account not make anything that an animal does that (merely) results in behavioral change a mental action. But it is hard to see how to do this on causalism without either

overintellectualizing nonhuman animals or simply positing an otherwise arbitrary difference between human mental representation and nonhuman animal representation.

The powers framework can resolve the issue here at what I think is a minor expense, namely, denying that nonhuman animals perform mental actions proper. On the powers framework, epistemic creditworthiness is indicative of something mental, where "mental" applies to organisms capable of forming and using truth-conditional representations. Epistemic creditworthiness is indicative of something mental because the distinguishing feature of human mindedness consists both in the capacities to construct and use such representations and in corresponding the discursive social praxes upon which the acquisition and mastery of these capacities depends. Living by thought is definitional of our life form. And epistemic achievements are improvements in how we individually or collectively engage in our specifically *rational* form of life. The paradigm cases, then, are not those of successful *covert* epistemic achievements. They are cases of successful veridical representation of facts that result in changes to networks of discursive solution systems such that application of those solution systems would constitute (more) accurate or veridical representations. And such successful acts of representing need not be covert. Indeed, to acquire the capacities to produce them at all, one must be shown overt examples of the acts of representing that are possible for an agent with the improved network and must practice overt performances of such acts prior to one's own network being thus improved. Once one has acquired these capacities, one can and often does go on engaging in overt exercises of them in instances of explicit learning.

On the powers framework, then, when Otto and Inga remember the location of MoMA overtly and covertly, respectively, they are connecting (or strengthening the connection between) a discursive solution system for representing the location of MoMA and other discursive solution systems that make us of its location in further discursive representations. Moreover, they are connecting (or strengthening the connection between) the former solution systems and at least one non-discursive solution system, namely, the one for going there. That Otto is extending his remembering through the use of his notebook marks no significant difference between him and Inga with respect to what both are doing. That said, we are tempted to credit other nonhuman animals with learning, exploration, discovery, solving problems, figuring things out, and so on because, in doing whatever constitutes these achievements, these organisms are shaping their own agency through the discovery, generation, or exploitation of novel opportunities for action. And the powers framework allows one to say that these are genuine *practical*-epistemic achievements. For they are changes to the animal's network of solution systems such that the application of those solution systems would constitute (more)

efficacious interaction with the world. But they are not the sort of achievements whose warranting epistemic credit draws one to distinguish mental and non-mental actions. Those achievements are reflected in changes in how human beings individually or collectively discursively represent and are assessed in terms of whether those changes trend towards (more) accurate or veridical discursive representations. And it is in recognition of such achievements that we ought to identify them with setting the standard and scope of mental action. In which case, mental actions are uniquely human. Moreover, mental actions not only can be extended but often are extended and acquisition of the relevant intentional action concepts requires some amount of extended performance and practice. That said, nonhuman animal learning, discovery, exploration, and the like are likewise achievements that set a standard and scope for a different kind of action that might well be collectively called "animal cognition." Moreover, none of the foregoing is meant to suggest that human mindedness is somehow more important or better than nonhuman animal mindedness. It is only meant to gesture at differences in life form that warrant care in how we approach correspondingly distinct kinds of mindedness.

2.4 Objections and Replies

In this section, I consider two objections. First, one might object to the very idea of mental action beyond a motley of "preparatory, ground-setting, tuning, retuning, shepherding" acts, acts of "*concertion*" or "dragoon[ing]" of thoughts or wandering minds (Strawson 2003: 231-232). Second, one might object to the powers framework's rejection of the analysis of intentions according to which they are propositional directive attitudes that are realized in the activity of the agent's neurocognitive architecture and whose contents stand in or go proxy for some state of affairs that she is to bring about. Without such attitudes, the objection goes, there is nothing to distinguish cases of lucky success from genuine intentional action because there is nothing to control or guide the relevant behaviors (Wu, personal correspondence). I unpack and respond to each objection in turn.

2.4.1 Strawson's Dilemma

I take Strawson (2003) to be exemplary of the first objection. Strawson raises the worry that there are no mental actions outside of those involved in setting the stage either for the introduction of content into consciousness or the direct manipulation of such content. The worry is expressed in terms of a dichotomy between control and automaticity. Strawson does not give a definition of either control or automaticity. But one can surmise what he has in mind from his contrasting the two. Automaticity is a matter of an event or process coming about without

the agent's initiation, intervention, or guidance and in such a way that the agent at least has a very hard time preventing the occurrence or stopping it midway (2003: 228*ff.*). Possessing control, then, is a matter of an agent's being at least in a position to initiate, intervene on, otherwise guide, and stop the relevant event or process. And exercising control is a matter of the agent's initiating, intervening on, guiding, or stopping the occurrence of the event or the completion of a process over which she possesses control in accordance with her intention(s). Both possession and exercise of control require that the agent have access to and capacities for interacting with whatever mechanism are involved in the occurrence of the relevant event or process.[68] We presumably have such access and capacities when it comes to the (at least gross) mechanisms of bodily action, and this explains why we are capable of intentional bodily action. And we have such access and capacities when it comes to the mechanisms involved in setting the stage for content-involving mental events and processes: we can initiate, intervene on, otherwise guide, or stop instances of saying or repeating things to ourselves, taking the results of a bit of mental activity and piling them up to be gone through one by one, "actively receptive[ly] blanking" the mind so that missing or new contents have "a chance to arise," and directing attention (Strawson 2003: 232). All of these can but need not be mental actions because we possess control over their occurrence and sometimes exercise that control. But the same is supposedly not true when it comes to mental events or processes that involve the introduction of content into consciousness or the direct manipulation of such content. So, for Strawson, for an occurrence to be an intentional action in general, it must be controlled and not automatic. And for an event-type to count as a type of action, an agent must at least sometimes possess control over its tokens. But because the mental events and processes that involve the introduction of content into consciousness and those that involve the direct manipulation of such are *always* automatic, they cannot be controlled. So, there are no mental actions outside of those involved in preparing the way for such events or processes. Such events or processes are merely "a matter of ballistics, mental ballistics" (241).

As I understand it, Strawson's main argument takes the form of a dilemma (234-238). It goes as follows: suppose for *reductio* that an agent can intentionally bring into consciousness or directly self-consciously modify some specific content, *p*, through thinking, remembering, imagining, judging etc. that *p* as a result of

[68] For an account of automaticity and control that is strikingly similar to Strawson's in cognitive science, see Schneider and Shiffrin (1977). Wu (2013) holds onto the idea that control implies a lack of automaticity and *vice versa* relative to the various aspects or properties of an event or process rather than relative to the occurrence of an event or process itself. But several philosophers have denied that automaticity implies a lack of control (Fridland 2015, 2017a; Arango-Muñoz and Bermúdez 2018; Douskos 2019). I am sympathetic to denying the dichotomy. However, in early work (2022), I accepted it in order to show that, even then, Strawson's argument does not show that episodic remembering is not at least sometimes an action (and a skillful one at that). I now think that argument was unsuccessful.

acquiring an intention to think, remember, imagine, judge, etc. that *p*. Either *p* is contained in the relevant intention or it is not. Starting with the first horn, if the intention contains *p* then any further mental action is made superfluous once the intention is formed. If the agent is consciously aware of the intention then she is already conscious of *p* by virtue of its being contained in the intention. Then the question is whether the formation of the intention is itself something the agent brings about by a second preceding intention or not. If it is insisted that the agent brings about a second preceding intention intentionally, a vicious regress of intended intentions opens up. If it is allowed that the formation of the intention is not the agent's doing then neither is *p*'s coming into conscious. The latter option is obviously preferable to the vicious regress. If the agent is somehow *not* consciously aware of the intention then bringing the content of the intention into consciousness is simply a matter of directing her attention to that intention. She need not engage in any further mental action than that. And Strawson allows that directing attention is a mental action. Either way, the first horn of the dilemma leads to a dead-end: the agent cannot intentionally bring *p* into consciousness or directly manipulate *p* in consciousness through thinking, remembering, imagining, judging etc. that *p* as a result of acquiring an intention to think, remember, imagine, judge, etc. that *p*.

Turning to the second horn of the dilemma, if *p* is *not* contained in the intention then the intention is to bring into consciousness or directly self-consciously manipulate some content (plausibly) related to *p*. One might then say to oneself "think about *p*," "remember *p*," "imagine *p*," etc. Or one might engage in "silent mental imaging of words (as sounds or visual marks, say)" (235). Yet, these "action[s] of setting oneself to produce some content or other" is not itself the bringing of *p* into consciousness nor the direct self-conscious manipulation of that content. The relevant intention, then, only motivates and (non-deviantly) causes or guides whatever stage-setting or preparatory acts the agent can muster to trigger the bringing of *p* into consciousness or the direct manipulation of *p* in consciousness. And the process of thinking, remembering, imagining, judging, etc. so triggered is not subject to any further guidance by the agent. It is therefore not something the agent herself does.

The proponent of mental action, then, impales herself no matter which horn she chooses. That is, in neither case is the desired content's coming into consciousness or the direct self-conscious manipulation of that content an action. At most, we intentionally set the stage to, say, carry out a bit of mental arithmetic, to imagine or remember "*p*," to reason about whether "*p*," to deliberate about whether to do *A*, and so on. We cannot intentionally perform such conscious, contentful mental actions themselves.

There are at least five ways of responding to Strawson's dilemma. First, one can point out that Strawson's identification of mental action with stage-setting or preparatory acts misses some other mental actions that do not function either to bring some particular content into consciousness or to directly self-consciously manipulate that content (for instance, see Mele 2009 and Brent 2023). Maintaining contents already in consciousness can be an agent's doing as well. Meditation and maintaining perceptual contact with something by attending to that thing would count as two kinds of mental action that Strawson does not seem to count but which, by his own lights, he should (on the former, see Brent 2023; on the latter, see Crowther 2009 and Wu 2023). Second, one can insist that the relevant intention neither contains the exact content one intends to bring about nor is an intention to engage in stage-setting or preparatory acts. Rather, the intention functions to specify a content that the agent brings about in performing the corresponding mental action. For instance, the intention to remember what one had for breakfast yesterday specifies a content to be recollected, say, "oatmeal," which intention then (non-deviantly) initiates and guides mnemic processing towards recollecting having oatmeal for breakfast (for versions of this response, see my 2022). The relevant content of the intention, in this case "remember what I had for breakfast yesterday," co-refers to what is represented in the resulting action, in this case the forming of the memory of having oatmeal, such that the content of the latter answers to the content of the former (Peacocke 2021, 2023).[69] Third, one can question whether Strawson's list of stage-setting or preparatory mental actions are not themselves acts of bringing content into consciousness or directly self-consciously manipulating such content. In which case, it is unclear why thinking, judging, remembering, imagining, and so on do not also count as possibly under our control. Indeed, once we adopt Strawson's control-automaticity dichotomy and once we accept his claim that these mental events or processes are not intentional actions because automatic, it is unclear why *bodily* movement does not also end up on his figurative hitlist. After all, all bodily movement involves some amount of automaticity. We could not move our bodies at all if this were not the case. So, if we adopt Strawson's dichotomy and accept that thinking, judging, remembering, imagining, and so on are not intentional actions because automatic then we should think that there are no intentional actions whatsoever (for a version of this response, see Wu 2013). But such a result is absurd. So, we should reject Strawson's claim that mental events or processes of bringing contents into

[69] One plausible way to understand the relationship between the content of the intention and that of the mental action is as having the same reference but distinct Fregean senses. Thanks to Bob Brandom and Madeleine Levac for suggesting this point.

consciousness or directly self-consciously manipulating that content cannot be actions. I will get to the fourth and fifth responses in assessing the first three.

The first response strikes me as not really an objection to Strawson so much as a friendly emendation. On a close reading of him, he seems to at least implicitly acknowledge such acts as meditating or perceptually attending to something as mental actions. The second, while on the right track, now strikes me as insufficient. Sure, the content of the agent's intention to remember what she had for breakfast yesterday co-refers to "oatmeal" if the agent had oatmeal for breakfast yesterday and if the mnemic processing that occurs as a (non-deviant) effect of her having that intention produces the appropriate mnemic content. But the connection between what the agent wants and what she does to achieve, secure, or get it still seems too tenuous. For the second response only allows that the agent enjoys some amount of *indirect* control over her cognitive processes. She forms the intention, e.g., to remember what she had for breakfast and then "wait[s] for [the] content to come" (Strawson 2003: 237). And such indirect control is not enough to say that she really has control over her remembering. Rather, on the second response, she has as much control over her thinking, remembering, imagining, judging, and so on as she does over, say, her tolerance for spiciness on specific occasions. An agent can build or lose a tolerance to spiciness. She can intend to increase or decrease her tolerance and this intention can lead to repeated acts of eating increasingly spicy things or abstaining from eating those things. And she can intend on some occasion, say, to weather the heat of eating a ghost pepper. But this does not mean that she really has control over her tolerance itself if what having control in such a case amounts to is that the agent is in a position to simply dictate the amount of spiciness she can handle on *this* occasion, to intervene on her experience of spiciness in the moment to increase or decrease "the heat" she experiences, or to altogether stop the relevant sensation at a whim. Assuming Strawson's control-automaticity dichotomy, then, how she experiences eating spicy things on specific occasions is not an intentional action of hers. The problem with the second response to Strawson's dilemma is that it does not by itself show that the agent has such control over her thinking, remembering, imagining, judging, etc. either. The intention to remember what she had for breakfast yesterday bears the same relationship to her remembering that she had oatmeal that her intention to tolerate the ghost pepper does to her tolerating it. Saying that the agent's intention co-refers to the contents produced by her cognitive processes does not go any further to showing that she has more control over them than she does over her tolerance for spiciness on specific occasions. So, the second response does not by itself show that there are mental actions of the sort that Strawson denies.

The third response is a genuine objection to Strawson from within the causalist camp. It accepts Strawson's control-automaticity dichotomy and points out that the relationship between intentions and the actions they (non-deviantly) cause and guide is pervaded by automaticity. Assuming causalism, we have two options: accept that there are no intentional actions whatsoever or allow that there are mental actions of thinking, judging, remembering, imagining, and so on. Wu's (2013) argument in particular is presented as a *reductio*. Accepting that there are no intentional actions whatsoever is intolerable because it flies in the face of common sense and is inconsistent with Strawson's own thinking. So, we should take the second option and accept that there are mental actions of the sort that Strawson denies.

So far, so good. But there are two further options: reject Strawson's control-automaticity dichotomy or reject causalism altogether. The fourth response to Strawson's dilemma, then, is to reject his dichotomy between automaticity and control. At least some (fully) automatic processes can be controlled such that their coming about are nonetheless intentional actions. For instance, suppose a paraplegic agent sets up a machine that strikes her knee with a reflex hammer whenever she presses a button. She forms an intention to kick, and this intention (non-deviantly) causes and guides her to press the button, the machine does its thing such that the reflex hammer strikes her knee, and she undergoes a reflex kick. Assuming that the reflex kick is fully automatic, it is nonetheless a behavior that is (non-deviantly) caused by her intention to kick.[70] So, on a form of standard causalism that rejects Strawson's control-automaticity dichotomy, the agent's reflex kick is a non-basic intentional action of hers despite being (fully) automatic. If this is right then even if we accept that cognitive processes are (fully) automatic, this by itself does not show that such processes are not also at least sometimes controlled and, in those instances, (non-basic) intentional actions. Strawson himself even entertains this response (236). And the second response to Strawson's dilemma can provide the cases that the fourth response would urge us to call (non-basic) intentional actions, namely, those in which the content of the intention in

[70] Colin Allen has informed me (though I cannot remember when) that reflex kicks are not actually fully automatic. That is, we can modulate reflex kicks if we try. We might even be able to prevent them. And there is some empirical research suggesting that we can voluntarily suppress stretch and cough reflexes (for instance, see Ludvig et al. 2007; Hegland et al. 2012). I use the example of reflex kicks partly because Strawson contrasts the voluntary both with the involuntary and with reflexes. I also use the example partly because accepting that reflexes in general are not really fully automatic leaves the domain of automatic human *behavior* quite sparse. We do not have many fixed action patterns. And, arguably, (fully) automatic processes like digestion do not count as behaviors at all. All of this is to say that the control-automaticity dichotomy starts to look less realistic and increasingly caricaturesque once one starts considering the finer details of the dynamics between organism and environment. The rejection of the dichotomy I am presenting here assumes its initial intelligibility. But one could reject it by pointing out that the dichotomy has much less empirical backing than it might at first seem.

question co-refers to the product(s) of the relevant cognitive processes and (non-deviantly) causes and guides those processes despite their occurrence being (fully) automatic.

Finally, the fifth response is to reject causalism altogether. Strawson and proponents of the first through fourth responses all accept some form of causalism. I have said what I find problematic about the first and second responses. The third and fourth responses strike me as problematic because they mislocate the problem. The real problem with Strawson's dilemma is that it misconstrues the agent's connection to her action. That misconstrual is presented most forcefully with the control-automaticity dichotomy. But it is brought out as well with Strawson's upstream assumption of causalism. That is, he assumes that for an event or process to count as an intentional action, that event or process has to be (non-deviantly) caused and guided by the relevant intention. The third and fourth responses accept the upstream assumption. The fourth response tries to show that some events or processes are such that the agent can bring them about (non-deviantly) through her having the relevant intention without her necessarily having (direct) control over those events or processes. The third response tries to show that tying intentional action to intention through the possession or exercise of control over the occurrence of an event or process *full-stop* results in absurdity.

But I think the fundamental problem lies with construing intentional action as behavior that is (non-deviantly) caused and guided by a truth-conditional propositional directive attitude. For on such a construal, agents relate to all of their actions at most as artifactual products of those directive attitudes. That is, on such a construal, a wedge is driven between the events or processes constitutive of action and the agent's directive attitudes through the introduction of causal relations that the agent herself cannot possibly relate to as agent. The agent is supposed to relate to her directive attitudes directly as their possessor—their formation just is her coming to have them. And their being hers is supposed to be part of what explains her going on to act. By virtue merely of having them, the agent is supposed to be motivated to act. The other part of the explanation is the (non-deviant) causal relations the attitudes are supposed to bear to the behaviors constitutive of her actions. In the case of non-mental action that extends beyond the body, the (non-deviant) causal relation is supposed to consist in enormously complicated activity in the neurocognitive architecture, the musculature of the body, and whatever mechanisms are involved in the things with which the agent interacts. In cases of non-mental bodily action, the (non-deviant) causal relation is supposed to consist in enormously complicated activity in the neurocognitive architecture and the musculature of the body. And in cases of covert mental action, the (non-deviant) causal relation is supposed to consist entirely in enormously complicated activity within the neurocognitive architecture. We can describe these relations in shorthand with talk of, say, sending an email,

walking, or thinking. But, for the reductive causalist, what makes the relevant happenings intentional actions is the instantiation of the relevant (non-deviant) causal relations themselves by virtue of the agent's having the relevant attitudes. And the agent cannot possibly relate to these relations as she is supposed to relate to the directive attitudes that are supposed to (non-deviantly) initiate such enormously complicated activity.

In ordinary cases of intentional action, the agent cannot relate to the activity of her neurocognitive architecture, the musculature of her body, or to the chaos of things going on in the world under purely physical descriptions *as agent*. None of them are, for her, things she can do. That is, none of them are the values of the variable "(I am) doing *A* (vis-á-vis *this*)" nor a "*this* that I can do." But her actions are! Yet, if reductive causalism is true, then the agent's relating to her actions as agent is only possible by dint of the shorthand descriptions happening to match the effects (non-deviantly) brought about through the instantiation of the relevant causal relations. There is nothing essentially tying what the agent does to her intention as falling under such shorthand descriptions except that such effects might regularly but contingently come about as a (non-deviant) result of her tokening the relevant intention. She cannot know them as bits of self-knowledge, that is, as falling under true "*unmediated* conceptions[/thoughts] [...] of states, motions, etc. of this object here, about which I can find out (if I do not know it) [through mediated conceptions/thoughts] that it is [Elizabeth Anscombe]" (Anscombe 1975: 62; emphasis added; see also Schwenkler 2019: chapter 2*ff*.). Mediation is typical of our conceptions and thoughts concerning objects other than ourselves because such conceptions and thoughts come by way of perception or inference.

That said, we can find out things about ourselves the way we find out about objects other than ourselves. And we can find out things about our actions the way we find out about objects other than ourselves. In fact, learning about ourselves and our actions through mediated conceptions and thoughts is integral to our development as persons, moral agents and patients, citizens, and so on. But it is not how we know our own actions in the first instance. Yet, if reductive causalism is true then such mediation is required to know our own actions at all. For the agent's access to the causal relations whose instantiations by her intentions (non-deviantly) bring about the effects constitutive of her intentional actions could only ever be mediated, say, by the use of fMRI or by observation of "the material [she] is working on" (Anscombe 1957/2000: §48). From where the causalist is standing, the fact that we are the authors of our action only makes our knowledge of them on a par with knowledge of artifacts, that is, of objects other than ourselves that we likewise author or produce (for an endorsement of such a view, see Envine 2016, especially chapter 7).

By contrast, the powers framework would understand knowledge of artifacts that an agent creates or uses, on the one hand, as consisting in practical knowledge of the act of creation or use involving the materials or the artifact itself, respectively, and, on the other hand, as consisting in speculative knowledge of the material worked on or the topic of her intention. Much human action involves the creation or use of artifacts. They are not alien to our agency in the sense that rational agency first comes without them and then incorporates them. Rather, they are alien to our agency by being independent existences outside of us, whereas our actions are not independent existences outside of us. On the powers framework at least, intentional actions are expressions of our minds, extensions of us through their being instances of our interacting with the world. We are authors of our actions in the same sense that we are authors of ourselves. In taking action, we contribute to self-making and through self-making we shape our future action. We are not authors of our action in the sense that we construct our actions at each instance. Artifacts are things with which we interact and through use of which we interact with the world. Our own actions are not things with which we interact, and we do not use them so as to interact with the world. They just are our interacting with the world. Our actions do not have the kind of independent existence that artifacts do. And it is because of this that we know actions differently from artifacts, namely, as self-expression and *not* as objects that we produce.

The real lesson of the third response to Strawson's dilemma, then, is that looking for a way for the agent to relate *as agent* to the causal relations that her truth-conditional propositional directive attitudes are supposed to (non-deviantly) initiate in notions like guidance or control is doomed to failure. The *reductio* Wu (2013) presents is not that of Strawson's argument alone but of causalism in general. That is, we should accept that if reductive causalism were true then there'd be no intentional action whatsoever. Thus, we should reject reductive causalism.[71]

2.4.2 Wu's Worry

Wayne Wu, in conversation, voiced a worry about the powers framework's account of mental action. At the time of conversation, I had no idea how to articulate the view at which I was aiming. But I was convinced of the idea that at least some mental actions were habitual and that their being habitual meant that their performance

[71] Agent causalism that attempts to rest easy with the claim that the agent, in the capacity of a substance, is a cause is no better off (for example, see Chisholm 1964). For to say this is simply to replace the truth-conditional propositional directive attitudes of the reductive causalist with the notion of a substance-as-cause. It does not explain how the agent relates to what she (non-deviantly) causes as her action because it does not explain how the agent relates as agent to the causal relation between herself in the capacity of substance and her action.

did not require an attendant truth-conditional propositional directive attitude to (non-deviantly) cause or guide them. Since then, my conviction has evolved into something of a radicalization: *no* intentional action, mental or otherwise, habitual or otherwise, requires such an attitude. At least, not if the possession of such an attitude is understood to mean that there is a truth-conditional representation of what is to be done as to be done by the agent that is realized somewhere in her neurocognitive architecture and whose function is to (non-deviantly) cause and guide her movement such that her movement satisfies its content. Wu's worry was this: what marks the difference between a habitual mental action like figuring out the amount to tip at the end of a meal and a happy accident of the right amount simply "popping" into one's head? Given my recent radicalization, the urgency of Wu's worry only increases. For now, the question is this: what marks the difference between any intentional action and the relevant type of event occurring purely by accident? The reductive causalist has a ready answer: it is the presence and (non-deviant) causal efficacy of an intention, understood as a truth-conditional propositional directive attitude with the right content that is realized in the activity of the agent's neurocognitive architecture. By rejecting that there are such internally realized attitudes, the powers framework seems to be at a loss in marking the difference. Hence the worry.

In response, the worry requires making an illicit move. The move is to assume that if the powers framework does not locate the difference between intentional action and lucky accident within the causal order of happenings within the agent's neurocognitive architecture then the framework leaves no space either for such happenings or causation to play a role in its account of intentional action. But this is wrong. For the powers framework understands the embodied procedural knowledge-how constitutive of an agent's skills as *productive* knowledge. Such knowledge just is the agent's causal powers to instantiate the relevant means-end orders when and where appropriate. In which case, the framework makes room for causation both within the neurocognitive architecture and between that architecture and things in the world to play a role in its explanation of intentional action. The mistake of the worry and of causalism more generally is to identify that role with the tokening of truth-conditional propositional directive attitudes realized in the activity of the agent's neurocognitive architecture. There is an alternative available. Namely, the relationship between the *a priori* moments of activation, ongoing performance, and completion of an action. Recall from chapter 1 that, on the powers framework, that relationship is fundamentally teleological and, thus, normative. And it is grounded in the agent's practically cognizing the means to her ends. In the simplest case, she has a need in some circumstance and there is some target with which interaction of the appropriate type would satisfy that need in that circumstance. The obtaining of both these need- and target-facts and the agent's coming into (perceptual)

recognitional contact with them potentiates the appropriate type of interaction on her part (or generates an intentional action acorn for the appropriate action). That potential is actualized in her interacting with the target in the appropriate way. And that actualization is successful by culminating in the satisfaction of her need. The process just described is, again, a teleological unity consisting of these *a priori* moments that develops in characteristic stages along said moments. It is a formal description of any intentional action. The formal nature of these moments is what makes them *a priori*, their connection to each other as parts of an ongoing process is what makes them unified, and their being grounded in the agent's cogent practical cognition is what makes them purposeful or subject to appropriateness conditions.

In the simplest case of non-mental action, the mental representation involved in the teleologically unified process of an intentional action is a *de agendo* representation of the target as affording a certain type of interaction whose generation by the agent is appropriate for satisfying the relevant need. That *de agendo* representation consists in the potentiation of that type of interaction on the part of the agent. In human beings, there is dual potentiation: potentiation of such an interaction on the one hand and potentiation of a myriad of ways of discursively representing the target and the agent's potentiated intentional action on the other. The latter are activations of solution systems for discursively representing the target constitutive of recognition and for deliberation, expression of intention, or action description, explanation, or assessment, respectively. All such potentiation is realized in the activity of her neurocognitive architecture. What the powers framework denies is that such realization is constitutive of tokening a truth-conditional propositional directive attitude. In the simplest case of mental action, there is perception or non-perceptual *de agendo* representation either of a discursive representational vehicle as affording interactions constitutive of constructing and using a discursive representation so as to mediately practically access some further target in ways appropriate to satisfy some, presumably intellectual need. In cases of overt mental action, such perception, understood as potentiation, naturally develops into actualization in the form of the agent's using the relevant vehicle or constructing and using such a vehicle. In cases of covert mental action, such perception or non-perceptual *de agendo* representation itself directly potentiates a set of further interactions with the discursive representation that an overt mental action *would* produce. This latter potentiation is a non-perceptual *de agendo* representation of the absence as affording further opportunities to mediately practically access the target of the would-be discursive representation. In other words, this latter potentiation constitutes the agent's recognition of the content produced by a covert mental action and the transition between the former potentiation and the latter constitutes the performance of that covert mental action.

The latter development in cases of covert mental action is likewise realized in the activity of the agent's neurocognitive architecture. But, again, such realization is not the tokening of a truth-conditional propositional directive attitude. Rather, it is the potentiation of a set of related discursive activities directly through the potentiation of some one discursive activity for the construction and use of the relevant discursive representation. The potentiation of one discursive activity leading directly to the potentiation of a set of related discursive activities for mediately practically accessing the same target is a single transition in thought, a single "movement of mind." But series of such potentiations can occur, leading to whole trains of thought. On the other hand, if the appropriate set of related discursive solution systems are all activated upon the agent's perceiving an object, that potentiation constitutes her recognition of that object. But she is not thereby performing any mental action. After all, such recognition is not mediated by the potentiation of an act of discursively representing the relevant target.

The point for now is that the powers framework can allow in its explanations of both mental and non-mental action a role for causation and for happenings within the neurocognitive architecture and between that architecture and happenings in the world. And it can allow for such a role without appealing to the tokening of truth-conditional propositional directive attitudes realized somewhere in the activity of the agent's neurocognitive architecture. On the powers framework, the role such causation and happenings plays is effectively that of the matter of corresponding actions. That is, the activity of the agent's neurocognitive architecture and the causal relations in which it stands to other such activity and to happenings in the world stands to the whole-agent-level activity of interacting appropriately with some target so as to satisfy some need as matter to form. It underwrites such agent-level activity but is not identical to it. Such matter is initially given form by the structuring of the agent's causal powers into means-end orders through the rational transmission of skill. It is given form on particular occasions by her (cogently) practically cognizing the means to her end on those occasions. Solving a math problem by a bit of mental arithmetic, then, just is knowing *how* to go on using the quantity represented by the discursive representation that one *would* have constructed if one had performed a corresponding *overt* mental action. In a case such as figuring out the amount to tip after a meal, such knowing how to go on might well be writing down on the relevant line on the receipt the number one *would* have written down as a result of applying the relevant equation on a bit of scratch paper. And this is a matter of the agent's perceiving the line as affording an opportunity to write down that number. If the covert mental act of figuring out the amount to tip is a habitual one, the agent's perceiving the total cost of the meal activates an incorporating mental habit for applying the relevant equation, potentiating its application such that a set of

related discursive solution systems for interacting with discursive representations of that number is then directly potentiated. One of these latter systems is then very likely to be applied in an actualization of the relevant type of interaction in light of further potentiation constitutive of the agent's perceiving the tip line, namely, actually writing the number down on that line. Such covert mental actions are indeed underwritten by the activity of the agent's neurocognitive architecture. And her overt mental actions and non-mental actions are partially underwritten by that activity and by the causal relations it bears to happenings in the world. But such activity does not itself constitute the agent's representing, acting intentionally, or engaging in discursive activities. It is only in reference to the interaction that the agent generates or is primed to generate *as a whole agent* that such activity can be identified as underwriting that interaction or the potentiation thereof.

Returning, finally, to Wu's worry, consider the case in which the right number just so happens to "pop" into the agent's head as a matter of luck. On the powers framework, something like the following must be going on. The agent has a mental skill or incorporating mental habit for applying the relevant equation by constructing and using the relevant numerals and operators. Her perception of the total cost of the meal is the activation of that disposition. That is, her perception is the activation of her mental skill or mental habit, respectively. But, in this case, a set of related dispositions for constructing and using discursive representations of the relevant number just so happens to get activated upon coming into perceptual contact with the total. Such activation does not occur *as a result of* the activation of the mental skill or incorporating mental habit for applying the relevant equation. She just so happens to be prepared both to apply the equation and to write the right number down, to utter it, to ask if it is the right number, etc. That is, her being prepared to do these latter things is not wrought by a combination of practicing figuring out the amount to tip on previous occasions and cogently practically cognizing the application of the relevant equation to the total *on this occasion*. And because the activation of these latter dispositions is not the result of the activation of her mental skill or incorporating mental habit for figuring out the amount to tip, it is not her coming up with that number by means of a covert mental act of figuring out the amount to tip.

The right amount does "pop" into her head and does so accidentally because the activation of her dispositions for discursively representing that number constitutive of such "popping" is accidental. On the powers framework, both this activation and that of the agent's disposition for applying the relevant equation are activations of highly determinate parts of the shape *in concreto* both of her determinable skill for doing arithmetic and of her more determinable skill for doing math. The dispositions whose activation is constitutive of the right number "popping" into her head are even more determinate than the disposition for applying the

equation since the former are dispositions for discursively representing that particular number are more specific. What makes the number's "popping" into her head a lucky accident, then, is that the activation of her dispositions for discursively representing that number and the activation of her disposition for applying the equation are not connected by way of the latter's activation. Had they been so connected, that connection would constitute her cogently practically cognizing as a practicable means the application of the equation towards the end of figuring out the amount to tip such that she landed on the right number. And such cogent practical cognizing, in virtue of the *a priori* teleological unity it provides, would in turn constitute her intentionally figuring out the amount to tip either skillfully or habitually.

In fact, cases of accidental success in general are a matter of a lapse in the agent's practical cognizing that nonetheless results in what the agent was aiming to achieve.[72] Take a classic case of causal deviance (Frankfurt 1978): a first-time robber is at a party and is supposed to signal to his comrades the time to commence the robbery by means of spilling his drink. On classic Davidsonian belief-desire causalism, his having a desire to signal his comrades and his believing he can do so by means of spilling his drink come together in the formation of an intention to signal his comrades by means of spilling his drink. But his forming the intention unnerves him so much that he trembles, spilling the liquid accidentally. The robbery ensues. The intention causes him to signal his comrades by spilling his drink. But his doing so is not intentional. He does not spill his drink as a means to signal. Many formulations of causalism have been put forward to deal with such cases. What matters here is what the powers framework has to say about such cases, since they are also cases of accidental or "lucky" success. The short answer is that what happens is not in this instance a unity consisting in the agent's cogently practically cognizing the practicable means of spilling his drink to his end of signaling his comrades.

The long answer goes something like this: recall the notion of practical want from chapter 1. To practically want something an agent has to conceive that thing as in some way good and as being at a distance and has to be trying to get it. Her trying to get it, in turn, consists in her being engaged in movement or action that (she takes) is of use towards getting it and her taking it that it is gettable by that movement or action. The problem with our robber is that he does not really practically want to signal his comrades. He has a subordinate need to signal his comrades that falls under his need to reap the benefits of the robbery and there is a target with which he can interact and to which he therefore has unmediated practical access, namely, his drink. His perception of the drink as spillable and his subordinate need together activate an overlapping part of his

[72] I am here distinguishing between accidental success and non-accidental unwitting success. On the latter, see Appendices D and E.

discursive solution system for signaling and his solution system for handling middle-sized dry goods containing liquids. But he is not thereby engaged in an act of spilling his drink to signal his comrades. He conceives signaling his comrades as falling under some good and as being at a distance but is not yet trying to signal them. After all, he is panicking and this is getting in the way, impeding the "natural" development of the activation of the overlapping part of the relevant solution systems into ongoing performance. Such panic might be the product of his fear of getting caught, his having a competing need to not see himself as a criminal, his excitement at the prospect of such a big windfall, or any number of things. Whatever induces the panic, it drives activation of solution systems whose application is practically inconsistent with the application of the overlapping part of his discursive solution system for signaling and his solution system for handling items containing liquids. Which other solution systems are activated will be a function of whatever's causing him to panic. He is prepared to run away or celebrate prematurely at the same time as he is prepared to spill his drink. And the conflict between which of these mutually practically inconsistent solution systems is to develop from activation to application goes unresolved on this occasion. By hypothesis, he does none of these things intentionally. He trembles as a byproduct of the panic-induced conflict, not as part of his practically cognizing the means to his end. For, on the powers framework, he never gets the chance to affect the relevant unity. He only gets as far as taking signaling his comrades as falling under some good and as being at a distance. And while he might recognize his spilling his drink as a "*this* that I can do" he does not take that practicable means for lack of corresponding practical want.

In cases of accidental success where the agent *does* take some practicable means, she nonetheless does not affect the relevant unity. In such cases, she practically wants something but either her practical cognizing is not cogent or she makes a mistake in performance that she cannot then correct. Take another classic case of causal deviance: a would-be murderer aims a gun at her would-be victim and fires but misses. Yet, the sound of the gunshot causes a herd of nearby pigs to stampede, trampling her would-be victim to death. On the most straightforward reading of the case, the powers framework has it that the agent makes a mistake in performance that she cannot correct. She practically wants to kill her victim and cogently practically cognizes shooting him as her practicable means. But her performance falters. Before she can take aim again, the pigs have come and done her work for her. But had they not, she might well have landed her second shot, affecting the unity she was after. And even in the case as it is presented, she could say truly of herself then and afterwards "I was in the midst of killing him when…" The case can be embellished such that the agent's practical cognizing is not cogent: she believes falsely that her would-be victim has a bad heart and, therewith, that the sound of the gunshot will be

enough to induce a fatal heart attack. She fires into the air and is dismayed to see her would-be victim startled but otherwise perfectly healthy. Then the pigs come. In this version of the case, two type-(iii) facts upon which her success depends do not obtain, namely, that her would-be victim has a bad heart and that a good scare would induce a fatal heart attack. She cannot affect the unity she practically cognizes because those facts do not obtain. She miscalculates. On the powers framework, in neither of these cases nor in the case of the first-time robber do our agents do intentionally what just so happens to culminate in the achievement of their goals. And in all of these cases this is because these agents do not affect the unity that they are after. Such unities are instantiated by the agent through her cogent practical cognizing and her avoiding or correcting for mistakes in performance.

In addressing both Wu's worry and Strawson's dilemma, I have contrasted the powers framework with causalism. I also criticized causalism in my discussion of extended mental action. The contrast and criticism are connected: the causalist's mistake is to treat actions, on the one hand, as happenings at the level of neurophysiological and physical events and the causal relations between them and, on the other, as happenings that the agent relates to as agent. Committing to both ways of understanding actions is commendable but untenable. It is commendable because such dual treatment seeks to at once provide a strong empirical basis for understanding agency while nonetheless respecting its role in our lives. It is untenable because it requires holding together in one conception of action two fundamentally different kinds of conception of the same thing. So, the causalist vacillates between the conception of action as a neurophysiological and physical happening and the conception of action as an expression of a subject's agency.

There are exceptions though. Empirically minded causalists who feel the tension in attempting to give action this dual treatment have increasingly abandoned treating actions as happenings that the agent relates to as agent. Instead, they speak of biasing, goal-states in the executive system driving such biasing, and actions themselves as behavioral responses produced by such goal-state driven biasing. The resulting account of intentional action is unrecognizable as an account of the exercise of rational agency. Trying to save the account by positing an identity between truth-conditional propositional directive attitudes and goal-states will not do. For such a move misattributes to the activity of the agent's neurocognitive architecture something she achieves in being brought into the relevant discursive social praxes, namely, an ability to act on the basis of what she herself takes to be a reason. It thereby makes this achievement completely mysterious. By contrast, on the powers framework, an agent's ability to act on the basis of what she herself takes to be a reason is really an achievement of the relevant sort. She can only do this if her agency has been shaped through the rational

transmission of skills into embodied means-end orders that she can instantiate through her cogent practical cognition. And she can only do this if such shaping includes the fashioning of dispositions for discursively representing such orders as part of that rational transmission and practical cognizing. Such shaping and fashioning shapes her neurocognitive architecture, her musculature, her interactions with extra-bodily objects, and the relationships between these things so shaped. But the exercise of her rational agency in intentional action does not consist *merely* in their interaction in the absence of such shaping and self-shaping.

2.5 Conclusion

In this chapter, I laid out the powers framework's account of mental action. Mental actions are expressions of mental skills and mental habits. Mental skills are rational dispositions for the construction and use of discursive representations. Mental habits are volitional dispositions for such construction and use. And such dispositions are constituted by networks of calculatively structured intentional action concepts for such construction and use. Mental actions, then, are the activation or application of these networks. Overt mental actions are applications, that is, intentional actions of constructing and using in principle intersubjectively available discursive representations. Covert mental actions are activations, that is, the potentiation of intentional acts of such construction and use sufficient to potentiate further intentional acts of discursively representing the resulting content. I also sketched three upshots of the account. First, acquisition of any one mental skill brings with it the acquisition of a battery of related mental skills. Second, several mental skills can be incorporated into corresponding mental habits and subsequently subsumed under recognitional skills. Recognition, then, is a matter of being primed to perform a host of basic mental actions concerning the same object. Third, cogent practical cognizing pervades the potentiation of mental actions and the exercise of mental skills and habits. I likewise expanded on three dimensions of the account. First, human perception is recognition-based. Second, the powers framework's account of the unity of mental action nonetheless allows for type-individuating such actions by appeal to the type of problems that those actions function to solve and by the type of means by which they solve those problems. Third, overt mental actions are extended and extended mental action is the rule rather than the exception. Finally, I responded to two objections. In responding to them and in elaborating the dimension of the powers framework concerned with extended mental action, I hope to have put some pressure on at least reductive causalism both about mental action specifically and in general. Again, I do not take myself to have provided knock-down arguments against causalism in general or any particular form that it takes. Applying pressure is only meant to allow the powers framework's account of mental action some room to

breathe. In the following chapters, I turn my attention to imagination, providing a sketch of an account and then arguing that, given this sketch, episodic remembering turns out to be a sub-type of imagining.

3.0 Imagining as a Skillful Mental Action

3.1 Introduction

In a sense, we all know what imagining is. It is what kids do when engaged in pretend play, what an avid fiction reader does when engrossed in a gripping novel, what a cunning general does when taking up the perspective of their adversary, what a savvy interior designer does when considering whether a piece would fit the aesthetic of a room, what a philosopher does when working through a thought experiment, what a bored department head does when mind-wandering instead of reading this year's expense reports, and what we do when bringing images or thoughts of the (merely) possible before our mind's eye. The kids imagine the goings-on of the pretense, the fiction reader imagines the world of the novel, the general imagines thinking like their opponent, the interior designer imagines the arrangement of the room that includes the piece, the philosopher imagines the setup of the thought experiment, the department head imagines being anywhere but here, and we imagine whatever is before our mind's eye.

However, there is another sense in which we do not at all know what imagining is. It is unclear exactly what type of activity is common across the aforementioned cases. And it is unclear what features of imagining would allow it to contribute to such disparate behaviors as pretending, reasoning about possibility and necessity, mind-wandering, and so on. Indeed, even just these three behaviors appear to make conflicting demands on imagining: motivating action, justifying belief, and neither, respectively. For the appeal to imagination to do real explanatory work across the aforementioned cases, there must be some characterization of what it is that is both intelligible and explicates its role in contributing to the relevant behaviors. Unfortunately, attempts at such a characterization have met with difficulty. Such difficulties have led philosophers to doubt that imagining is a unified, potentially explanatory phenomenon after all. Indeed, after considering various kinds of imagining, Walton asks: "What is it to imagine? [...] should not we now spell out what [these kinds] have in common? Yes, if we can. But I can't" (1990: 19). And, similarly, Kind asks: "[I]s there such a thing as *the* phenomenon of imagining? [...] we must answer in the negative: There is no single "something" that can play all of the explanatory roles that have been assigned to [imagining]" (2013: 157; original emphasis).

In this chapter, I drill down on the powers framework's account of mental action as it applies to imagination. On this basis, I provide a novel, non-reductive, action-first, skill-based account of active imagining that can explain what is common to imagining across the behaviors it is invoked to explain. I call it the Skillful

Action Account of Imagining (the Skillful Action Account for short). According to this account, imagining is paradigmatically active. To imagine something actively is to (at least) potentiate the formation of a discursive representation of that thing, where such potentiated formation involves the agent's selecting the content of that would-be representation as a means to the performance of behaviors such as those listed in the preceding paragraphs.[73] I place such behaviors together under the heading of "imaginative project." So, for example, in a pretense game like 'the floor is lava', one player might respond to another's tossing a throw pillow onto the floor by potentiating the formation of a discursive representation of, say, an outcropping emerging from the lava as a means of carrying on the pretense. Completion of an imaginative project like carrying on a bit of pretense stands to potentiating the formation of the corresponding discursive representation as an end. The account thus bakes in the calculative, means-end order that the powers framework takes to be definitional of intentional action. Moreover, in the spirit of this conception of intentional action, I hold that a central constitutive feature of this order is the agent's exercising practical knowledge both of her act of imagining and of its having the relevant means-end order. In exercising that practical knowledge, our pretender knows *that* she is actively imagining the emergence of the outcropping and knows *why* she actively imagines this, namely, to carry on the pretense. Her knowledge in her intention to play constitutes her corresponding act of imagining. Both the means-end order and the agent's practical knowledge thereof are on the Skillful Action Account essential parts of what it is for her to imagine. Moreover, in taking active imagining to be the central imaginative kind, the account highlights the importance of both the means-end order and practical knowledge to understanding the nature of imagining.

The Skillful Action Account provides a characterization of imagining in terms of its paradigmatically being a type of skillful action. Human beings might well start out with a bare capacity to generate imagery in the absence of a corresponding stimulus at will.[74] But, even so, for this capacity to contribute to pretense and the like, agents have to learn how to imagine in ways appropriate to the completion of the corresponding imaginative project. And, as with other types of skillful action, this takes training and practice. So, although infants appear able to engage successfully in simple pretend play and in joint acts of pretense with non-infants (Nielson and Christie 2007; *cf.* Hess 2006), imagining appropriate to, say, completing a yearlong Dungeons & Dragons campaign requires an ability to imagine shaped by sufficient practice playing tabletop games. Similarly,

[73] Because most of the literature on imagination is concerned specifically with covert imagining, I will not mention actual formation of discursive representations going forward except where clarity dictates. I will likewise forgo distinguishing covert from overt imagining going forward except where clarity dictates. I leave open whether the creation of art or fiction count as overt acts of imagining on the powers framework, although I suspect that they do.

[74] Here I leave it open whether such imagery represents anything.

as with other types of skillful action, one's imagining can improve along certain dimensions and can be appropriate or not in relation to the type of imaginative project one is engaged in. For instance, a pro tennis player's serve is apt to be smoother and more uniform than an amateur's. And a pro tennis player's serve is apt to take different bodily movements as means and have more force overall compared to the serve of a pro badminton player.

The characterization of imagining as paradigmatically a type of skillful action provided by the Skillful Action Account can explicate imagining's role as a means in contributing to pretense, engagement with fiction, predicting or explaining others' behavior, reasoning about possibility and necessity, hypothetical and counterfactual reasoning about contingent matters of fact, mind-wandering, imagining for its own sake, and so on. According to the account, as with other types of skillful action, distinct acts of imagining can be put to use towards very different, even conflicting ends. Moreover, in their service as means, such acts can take on very different, even conflicting properties. For instance, painting as part of creating a work of art might well require the agent to do something that is very different or that even conflicts with what would be required of her in painting as part of renovating a home. Nonetheless, painting as a type of skillful action is identifiable despite this variation and conflict across its instances: it is the exemplification of some social praxis of at least minimally aesthetic expression through color. Skillful act-types tolerate variation and conflict in the role they play across their instances. If, as the Skillful Action Account has it, imagining is paradigmatically a type of skillful action then such variation and conflict across its instances is to be expected and is consistent with type-individuating imagining as a skillful mental action.

In what follows, I introduce the Skillful Action Account (Section 3.2) before applying it to explain, programmatically, the role of imagining across types of imaginative project (Section 3.3). I then show how the account elucidates the intentionality of imagining as part of its characterization of imagining (Section 3.4) and addresses passive instances of imagining (Section 3.5) before concluding by touching on six avenues for future research.

3.2 Advancing the Account

Here, in brief, is the account:

> **The Skillful Action Account of Imagining**
> An agent's potentiating the formation of some discursive representation is an instance of *covert* active imagining *iff*
>
> i) she is primed to construct that representation and to select its content such that they together constitute a means to her completing some project and;

ii) her completing that project consists in the exercise of her practical knowledge of her potentiated constructive and selective activity as a means.[75]

The account treats active imagining as the central imaginative kind and, thus, takes imagining in general to consist paradigmatically in the performance of this type of mental action. To simplify things, going forward, I restrict focus to active imagining among neurally intact individuals from adolescence through adulthood under normal conditions.

By "primed to construct that representation," I mean very roughly what Van Leeuwen (2013) calls "constructive imagining." That is, active imagining involves the agent's potentiating "coming up with [discursive] representations that have [such-and-such] content" by "combin[ing] elements of ideas from memory" in a way usually governed by background beliefs, conventions, and guiding principles specific to each type of imaginative project (221, 224).[76] (Recalling Section 1.3.1, the agent generates an intentional action acorn for constructing and using a discursive representation with the relevant content which she selects, drawing from memory). For example, an act of covertly imagining a dancing banana involves the combination of elements of the agent's conception of bananas with that of her conception of dancing. Elements of the former might include the typical shape, color, texture, and so on of bananas, while elements of the latter might include stereotyped movements of, say, freeform dancing. On the powers framework, the imagery involved in such imagining consists in the activation of perceptual recognitional skills where these skills are activated directly by the agent's potentiating the construction of a corresponding discursive representation whose content she selects.[77] I depart from Van Leeuwen in denying that (potentiated) construction on its own counts as a kind of imagining. For instance, the coming of perception-like imagery to mind unbidden might well involve the

[75] In cases of overt imagining, the formulation would be as follows. An agent's forming some discursive representation is an instance of active imagining *iff* i) her constructing that representation and selecting its content together constitute a means to her completing some project and ii) her completing that project consists in the exercise of her practical knowledge of her constructive and selective activity as a means.

[76] In keeping with the powers framework, such "ideas from memory" consist in networks of discursive solution systems for constructing and using discursive representations of whatever content the "idea" corresponds to. "Memory" in the relevant sense, then, amounts to the agent's being in possession of a conceptual repertoire and perceptual prototypes consisting in such interrelated solution systems.

[77] The activation of the relevant perceptual recognitional capacities is likely subserved by the endogenous activation of mechanisms involved in both perception and semantic processing. These processes might well occur unconsciously; but, to my knowledge, neither Van Leeuwen himself nor those who appeal to the notion of constructive imagining hold that its representational products are had unconsciously or without awareness. Thanks to Bob Brandom for pressing me to clarify this. On activation of perceptual mechanisms in imagination, see Kosslyn (1994, 2005), Van Leeuwen (2016), Langland-Hassan (2016), Winlove et al. (2018), Gauker (2020, 2021), Williams (2021). On activation of conceptual capacities in semantic processing in imagination, see Johnson-Laird (1983), Williamson (2016).

subject's combining elements of ideas from memory, however implicitly. But such an instance of construction is not thereby an instance of imagining, let alone an instance of active imagining.[78] To imagine actively, an agent must not only combine elements of ideas she has but must also select those elements in light of what she aims to do with the discursive representation that *would* result from actually constructing it.[79]

By "primed to select its content," I mean what Dorsch (2012) discusses under the heading of the "specific determination" of content (387-391). That is, active imagining involves the agent's "voluntary determination of which entities are represented as instantiating which properties" (389). In imagining a dancing banana, the agent does not just select the subject matter of her imagining. That is, she does not just hold an intention to imagine a dancing banana and then leave the rest to nature (*cf.* Strawson 2003). Rather, she exercises selective agency over which entity she imagines and its properties, namely, some individual banana that is moving itself in the relevant way(s). She might also potentiate the formation of a discursive representation of the banana as having a particular color, shape, texture, and so on. I likewise follow Dorsch (2012) in taking it that the agent's involvement both in potentiating the construction of the relevant representations as well as in fixing those representations' content distinguishes imagination from other mental faculties like perception and belief. The agent's ability to select what she imagines is part of the so-called "freedom of imagination." This freedom has been thought to be a distinguishing mark of the faculty since at least as far back as Hume (1777/2000, 1793/2007). I will have more to say about this aspect of the freedom of imagination in Section 3.4.

I want to pause for a moment on the contrast between imagining, on the one hand, and perception and inference on the other. When an agent imagines, at least some of what she imagines remains unfixed until she determines its content. By contrast, when an agent perceives or infers something, what she perceives or infers is in the good case fixed by the relations that obtain between the relevant targets and the (potentiated) mental skills or mental habits constitutive of such perception or (covert) inference. In perception, content is fixed by

[78] On the distinction between imagery and imagination, see Kind (2005, 2021). This claim applies as well to the formation of a discursive representation that one treats as non-actual. The bare potentiation of such a representation in the absence of aiming at completing any imaginative project, say, in cases of intrusive thought, is insufficient to count as imagining. As Walton puts it, imagining requires "doing something with a proposition one has in mind" (1990: 20).

[79] Here and below, in keeping with the powers framework, I will be using "formation" and "construction" to pick out two aspects of active imagining. Formation concerns the (potentiated) construction and use of a discursive representation, whereas construction is itself a part of formation. The (potentiation of the) formation of a discursive representation, then, already involves the selection of content by the agent sufficient for use in completion of the relevant imaginative project. The (potentiation of the) construction of a discursive representation, while not separable from selection of content in the case of imagining, does not by itself imply selection, since such (potentiated) construction can occur as part of other mental actions.

the relation of perceptual contact that obtains between the agent and the relevant target in her perceptible environment. Drawing from Springle's framework, that relation consists in the activation of the agent's skills or habits for generating certain types of interaction with the relevant particular(s) and the obtaining of the corresponding target-fact(s) and is mediated via some sensory medium. The connection established between the target and the agent through the sensory medium and the network(s) of intentional action concepts constitutive of the skills or habits that she has to that point developed together fix perceptual content. The agent, in coming into perceptual contact with some particular, is primed to generate types of interactions specific to the type of particular with which she is in perceptual contact. In human beings, such target-type-indexed interactions include constructing and using discursive representations that categorize the relevant particular, attribute properties to it, depict the relations that it stands in to other objects, and so on. Perceptual content, then, is fixed by a kind of "attunement" between the agent and her environment. In inference, content is fixed by the premises used. That is, as with the case of mental arithmetic discussed in the previous chapter, the content of the discursive representation that the agent infers is drawn from the content of the discursive representations she uses to infer it. That said, if an inference starts from a supposition, the content of that supposition is determined by the agent. So, her supposing something for the sake of argument is itself an imaginative mental action.

By "project," I mean goal-directed movement, whether bodily or mental, that can be pursued by means of imagining as well as such movement that must be so pursued. These projects include but are not limited to:

- imagining for its own sake (Dorsch 2012; Van Leeuwen 2013);
- pretense (Gendler 2000, 2003, 2006a,b, 2010; Nichols and Stitch 2000, 2003; Carruthers 2002, 2006; Nichols 2004; Picciuto and Carruthers 2014);
- engagement with fiction (Currie 1990; Walton 1990; Meskin and Weinberg 2003; Doggett and Egan 2007; Chasid 2019);
- predicting or explaining others' behavior (often called "mindreading") (Currie and Ravenscroft 2002; Heal 2003; Goldman 2006);
- reasoning about possibility and necessity (Gregory 2003, 2020; Williamson 2007; Kung 2010; Ichikawa and Jarvis 2012; *cf.* Fiocco 2007; Spaulding 2016);
- hypothetical and counterfactual reasoning about contingent matters of fact (Kind and Kung (eds.) 2016; Kind and Badura (eds.) 2021; Myers 2021a,b; Blomkvist 2022); and
- mind-wandering (Christoff et al. 2016; Irving and Glasser 2018; Irving et al. 2020).

Following the philosophical literature on imagination, I treat each listed behavior performed by means of imagining as constituting a distinct type of imaginative project. Completing imaginative projects consists in performing the relevant behavior by means of imagining. Importantly, on the Skillful Action Account, agents

must learn how to complete each type of imaginative project. There is no inbuilt faculty or module for pretense, engagement with fiction, mindreading, reasoning about possibility and necessity, and so on. If so, then active imagining is skillful in large part thanks to the agent's learning and mastering principles, methods, techniques, heuristics, and so on specific to the domain of each type of imaginative project through rational transmission of the corresponding skills in being brought into the corresponding social praxes.[80] Her being brought into social praxes of pretending, engrossed engagement with fictional narratives, figuring out what someone will do next, and so on quite literally shapes her capacity to imagine over time, making her a more refined and sophisticated imaginative agent.[81] I will have more to say about imaginative projects and imaginative skill in Section 3.3.

By "means," I have in mind a more immediate action by which some other action is performed or by which some end that is currently at a distance from the agent is achieved, secured, or gotten. Drawing from chapter 1, immediacy and being at a distance are interdependent formal notions that pick out relations among actions and ends. The means an agent takes are actions that suffice to close at least some of the distance between her and her goal. A sufficient means is an action that suffices to completely close such a distance. Recall our pretender playing 'the floor is lava'. In playing the game, her imagining that the floor is lava is immediate and is done in order to perform pretense actions like jumping from one outcropping to another, where, for instance, two sofas function as the appropriate props in the game. Unlike the real-world act of jumping between sofas, the pretense act of jumping between outcroppings is at a conceptual distance. The latter action's being at a conceptual distance is due to the fact that, in addition to performing the real-world act, the agent must recognize jumping between sofas as constituting a legitimate move in the game and must be motivated to pursue this move. The calculative relation between means and ends on display in this example can be iterated. For instance, the agent might perform the pretense act of jumping between outcroppings in order to carry on the pretense and might carry on the pretense, in turn, in order to entertain herself and her friends. Moreover, each means sufficient for closing at least some of the distance between the agent and her goal might consist in the performance of multiple actions. In the case under consideration, the agent both imagines and actually jumps in order to perform the pretense act of jumping between outcroppings. For other imaginative projects, for instance, engagement with fiction, the act of imagining, e.g., the goings-on of the world of the fiction, can itself be a sufficient means to their completion.

[80] Hereafter, I use "principles" to cover methods, techniques, and heuristics except where clarity dictates.

[81] On the idea of an agent's capacity for some action having a particular shape, see chapter 1, Locke (1974), and Small (2017). On the idea of an agent's capacity to imagine having a particular shape, see Hopkins (2023, 2024).

Finally, by "practical knowledge," drawing from chapter 1 (and Appendices A, C, D, and E), I mean the agent's knowledge of what she is doing such that what she is doing has the calculative means-end order that constitutes it. Practical knowledge is the agent's knowledge *in intention* (Small 2012: 135*ff.*; Campbell 2017: 17*ff.*). Our pretender knows that she is playing 'the floor is lava' as part of her intending to play and this gives her knowledge of the means that she is taking: her imagining, her jumping, her performing pretense actions, and so on.[82] The means-end order I have been drawing out of the example would not exist were it not for her practical knowledge. Acquisition of the agent's ability to practically cognize what she is doing such that she has and exercises practical knowledge of her doing it depends in part on her having the requisite knowledge-how. In the case under consideration, part of what allows the agent to play intentionally and, thus, knowingly is her knowing how to pretend as well as her knowing how to play 'the floor is lava'. Such knowledge-how is gained through practice. Yet, once our pretender has acquired this knowledge-how, she can readily apply it whenever

[82] Note that while I hold that sophisticated imaginers might apply the concept IMAGINE in giving voice to their practical knowledge, I allow that children might only describe themselves as, say, pretending that the floor is lava. I take this to be due to children not having yet mastered an additional discursive skill to deliberate about or express through discursive means the content of their other acquired skills. As Anscombe, citing Aristotle, puts it "one does not deliberate about an acquired skill; the description of what one is doing, which one completely understands, is at a distance from the details of one's movements, which one does not consider at all" (1957/2000: §30). I take it that this means that, like nonhuman animals, our most immediate actions can be transparent to the material that we're working on (see Section 1.3.1) or, in our case, to the actions that extend beyond those most immediate ones. Indeed, pro-athletes find it difficult to describe their bodily movements in detail. And I take it that only after some discussion (bordering on elenchus) would a sophisticated imaginer accept that they do not literally construct, say, a picture in their head when they *covertly* imagine. The suggestion here is, first, that children are already skilled at imagining and that adults are for the most part master imaginers but that, second, this does not mean that either children or adults are thereby in a position to fully articulate their imaginative acts, since these acts are almost never at a distance. In an earlier version of this chapter, I claimed that imaginative acts are unified by being explorations or navigations of modal space, by which I meant sets of possible or impossible worlds (for a logic on which this is the case, see Berto 2017, 2018, 2023). I have since removed this from the account proper on the grounds that even fully realized imagining agents cannot have practical knowledge of such exploration or navigation. Rather, I believe that such exploration or navigation amounts to what McDowell calls "a suitable happening," that is, "a happening required for there to be a performance that conforms to a description 'of a type to be formally the description of an executed intention' [(Anscombe 1957/2000: §48)], where the content of the description corresponds to the content of the knower's intention" (2022: 220-221). Such happenings include those picked out by descriptions that make up the "chaos" in which the order of intentional action can be found, say, in "handing two bits of paper to a girl" when paying a bill or "dragging metal objects about" when setting up a mortar (Anscombe 1957/2000: §43). My suggestion is that "exploring modal space" or "navigating modal space" as part of active imagining, where this means locating sets of possible worlds, are suitable happenings rather than what the imagining agent has practical knowledge of. And while descriptions like "being primed to construct a discursive representation and to select its content" might also go beyond what agent has practical knowledge of, I take it that such descriptions are much closer to what is, in all likelihood, the basic mental act of imagining than is "exploring modal space" or "navigating modal space." That is, I take it that the former descriptions are of a now non-isolable means for most who have acquired and mastered the skill of imagining and, so, do not pick out suitable happenings while the latter descriptions do. Thanks to Nick Wiltsher for pushing me to clarify this point.

the opportunity arises. Because practical knowledge stems from the agent's practical cognition, that is, from her intention, it is knowledge unlike that drawn from belief or perception in that she possesses it without drawing on evidence (Campbell 2017: 18-19). Our pretender imagines that the floor is lava ultimately because she possesses practical knowledge of her end in intending to play.[83] Such knowledge contains her knowledge of the means she takes to that end. In the case under consideration, the agent does not need to observe or infer that she is imagining or that her imagining is efficacious to know that she is imagining, what she is imagining, or why. She does not rely on observation or inference because her imagining itself stems from her practical knowledge of it in intending to pretend. I will have more to say about the role of practical cognition in imagining in Sections 3.3-3.4.

Recall our pretender playing 'the floor is lava'. According to the Skillful Action Account, she actively imagines that the floor is lava because, in intending to carry on the pretense, she potentiates the formation of a discursive representation of the floor's being lava as a means to that end. By contrast, consider an agent engaging in Tolkien's *Return of the King* and, in so doing, imagining the scene where Frodo Baggins and Samwise Gamgee enter Mount Doom. According to the Skillful Action Account, this agent actively imagines, say, that Frodo and Sam are standing at the precipice, looking down on a molten pit whose flames are roaring all around them. He is doing this because he potentiates the formation of a discursive representation of this event in Frodo's and Sam's journey as a means to engaging in the fiction. In the one case, the pretender's active imagining is her potentiating the construction of a discursive representation with properties that allow it to bear both a motivational and justificatory relation to certain pretense actions, where that potentiation is sufficient to further potentiate those pretense actions. In the other, the fiction reader's active imagining is his potentiating the construction of a discursive representation with properties that allow it to bear only a justificatory relation to his active and immersive reading of *Return of the King*, where that potentiation is sufficient to further potentiate that active and immersive reading.[84] In the latter case, such secondary potentiation can take the form

[83] That an action is constituted by the imposition of the corresponding means-end order and that this order is itself given form by the agent's practical cognition is what Anscombe and others mean when they claim with Aquinas that practical knowledge is "the cause of what it understands" (1957/2000: §48; see also *De Anima*: 433a29-30; *Summa Theologica:* I, q. 14, a. 16c; Falvey 2000; Moran 2004; Thompson 2008; Setiya 2008; Small 2012; Schwenkler 2015, 2019; Campbell 2017). I acknowledge that this claim is controversial in action theory (e.g. Shepherd and Carter 2023). A full defense of it would go beyond the scope of this chapter. However, see Appendix E and Small (2012) for full-throated defenses and Schwenkler (2019) for a defense that the claim holds ordinarily or for the most part.

[84] This is not to say that the content of the representation, so potentiated, which the fiction reader imagines cannot be made to motivate action at all. After all, he could go on to play a game of pretense wherein he is Frodo and his friend is Sam. But to do so is to move from one imaginative project, engaging in the fiction, to another, namely, pretense. And

of imagery or of further thoughts about the state or nature of the fictional world. And, in either case, on the powers framework, such imagery and further thought is the activation of recognitional skills. What remains the same between the cases of pretense and engagement with fiction is the means-end order specific to imagining and each agent's practical knowledge thereof. That order and their practical knowledge are thus essential parts of what it is for them to imagine.

In taking active imagining to be the central imaginative kind, the Skillful Action Account highlights the importance of both the means-end order and practical knowledge to understanding the nature of imagining itself. More than that, the account sets out the form specific to imagining as a type of mental action. By virtue of being a type of action that can be performed intentionally, imagining is a type of happening characterized by a means-end order that is definitional of it.[85] The Skillful Action Account explicates the means-end order specific to active imagining and, thus, characteristic of imagining in general. According to it, that order consists in, on the one hand, the agent's potentiating the construction of a discursive representation and her selection of its content as means and, on the other, her completing the corresponding imaginative project as end. In proposing this account, I thus adopt a novel position concerning the metaphysics of imagination. In effect, what I am proposing is an action-first non-reductive ontology of imagination (for a contrasting view, see Langland-Hassan 2020). My aim in making this proposal is in the first place programmatic: I want to provide a unifying account of imagining that has materials sufficient to explain its role in contributing to the completion of imaginative projects. It is to this that I now turn.[86]

---

making such a move requires changing the corresponding imaginative discursive representation, namely, by potentiating the formation of a discursive representation of himself as Frodo and his friend as Sam and by potentiating the formation of a discursive representation of their current environment as the mouth of Mount Doom. The point for now is that there is a core to the variety of uses to which we put the imagination and that this core consists in a particular type of mental action. In the following section, I touch on the differences that crop up between imaginative projects, how such differences make a difference to the agent's practical knowledge of what she is doing in imagining, and how this difference in her more specific practical knowledge makes a further difference to her active imagining while nonetheless leaving untouched her more generic practical knowledge of potentiating the construction of a discursive representation and selecting its content as a means of completing such-and-such imaginative project. Thanks to an anonymous reviewer for pushing me to clarify how distinct types of imaginative projects fit within the Skillful Action Account.

[85] Recalling chapters 1 and 2, this is because there being appropriate descriptions of these types of happenings in the first place depends on "our possessing the *form* of description of intentional actions" (Anscombe 1957:2000: §47; original emphasis). Considered at the level of action generally, this form is just that of the means-end order (1957/2000: §26).

[86] One might worry that the Skillful Action Account is not particularly novel, citing Peacocke (1985) and Noordhof (2002) as examples of accounts of (imagistic) imagining in which the notion of an imaginative project plays a central role (and I myself would add Dorsch 2012 to the list). What is novel about the Skillful Action Account is not that it invokes imaginative projects as central by providing ends which require the relevant kind of means. Rather, what is novel is that the action taken as a means to such ends is skillful, where this means that it is learned through training and practice, that such

3.3 Applying the Account Part I: The Role of Active Imagining in Imaginative Projects

Recall the problem with which I began. On the one hand, we seem to know what imagining is and to be able to identify it as playing a key role across a number of disparate types of imaginative project. It seems that imagining involves the agent's ability to potentiate the formation of certain kinds of discursive representation at will (*cf.* Strawson 2003). And it seems that this potentiated discursive representation formation activity is implicated in pretense, engagement with fiction, mindreading, reasoning about possibility and necessity, hypothetical and counterfactual reasoning about contingent matters of fact, mind-wandering, and so on. On the other hand, these behaviors seem to make conflicting demands on imagining such that it is difficult to see what is common to the activity across its contributions to them. We are thus hard pressed to provide a characterization of imagining that is both intelligible and can explicate its role across the relevant types of imaginative project. The problem, then, is to find a unity within this multiplicity. Call this the Unity-in-Multiplicity Problem.

The problem is especially acute once the functions ascribed to imagining are considered. Given what imagining is thought to do within each type of imaginative project and what features of imagining are thought to allow it to do this work, it can seem that the activity is attributed conflicting functions. For instance, in pretending, imagining the goings-on of the pretense is supposed to be such that it motivates action in the actual world. By contrast, in engaging with fiction, imagining the world of the fiction is supposed to be such that it does *not* motivate action in the actual world. In hypothetical and counterfactual reasoning about contingent matters of fact, imagining what might occur or what could have occurred is supposed to be such that it justifies beliefs closely tied to the actual world. In many other types of imaginative project, imagining is supposed to be such that it does *not* justify beliefs closely tied to the actual world. In mindreading, imagining being in another's shoes is supposed to be such that it both justifies beliefs concerning the actual world and motivates action in the

---

learning results in the acquisition of abilities whose manifestations are exemplifications of the imagining agent's knowledge-how to imagine in ways appropriate for the completion of the relevant projects, that such appropriateness stems from domain-specific principles grounded in social praxes for engaging in the relevant projects and embodied in the agent's practical knowledge, and that the proficient imaginative agent's exercise of the relevant abilities amounts to her exercise of that practical knowledge of her imagining such-and-such as a means to her, e.g., pretending, engaging with fiction, etc. While others, in particular Kind (2021), have noted the skillful nature of imagining, the Skillful Action Account and the present chapter aim to elaborate (programmatically) in what the skill in question consists. And while Dorsch (2012) provides what is, to my knowledge, the only other unifying account of imagination in terms of action, the Skillful Action Account differs from his ACT account by flipping the dependence relation between active imagining and imaginative projects. Where Dorsch claims that imaginative projects depend on a kind of mental action that itself counts as imaginative independently of any such projects in order for such projects to count as imaginative, I claim that such projects are, by virtue of what they demand of the agent, part of what makes the act of imagining imaginative. Thanks to the anonymous reviewer for pushing me to clarify these points.

actual world. In many other types of imaginative project, imagining is supposed to be such that it *neither* justifies beliefs concerning the actual world *nor* motivates action in the actual world. In engagement with fiction and (at least certain instances of) mindreading, imagining is supposed to be such that it gives rise to affective states. In reasoning about possibility and necessity, imagining what possibly is or what must be is supposed to be such that it does *not* to give rise to affective states. Given conflict in the functions ascribed to imagining and in the features of instances of imagining that are supposed to allow it to perform these functions across types of project, the trouble becomes saying what imagining is that tolerates this conflict and does not just resolve into a heap of accounts specific to each type of imaginative project.[87]

One response to the Unity-in-Multiplicity problem has been to give up altogether on finding a characterization that can do both, as Walton (1990) and Kind (2013) suggest. Kind puts the motivation for this response best:

> [T]he explanatory burden that imagination must carry varies greatly from context to context. Not only do features that are essential to imagination in one context drop out entirely in another context, but even worse, features of imagination that play an essential role in one context are sometimes inconsistent with features of imagination that play an essential role in another context. [...] When philosophers invoke imagination to explain one [type of imaginative project] [...] the thought is that there is something special about imagination itself that can do the explanatory work. In each individual context, this claim may well seem plausible. But once we look at the contexts together, the initial plausibility of the claim dissipates (2013: 157).

[87] One might worry that this characterization does not distinguish between variation in functional role and complexity of a function. After all, imagining might have the same functional role across its instances and may only have salient differences given differences in their inputs. If so, then the air of conflict disappears: all imagining functions, say, to construct or potentiate the construction of a discursive representation whose content is fixed by the agent. Instances of imagining can appear to differ from each other to the extent that varying inputs have distinct downstream effects on the performance of that function. In response, the Unity-in-Multiplicity problem goes deeper. In particular, what is at issue is not only that the inputs of instances of imagining vary but the features of those instances and their outputs seem to vary as well. For example, in a game of pretense like 'the floor is lava', the (potentiated) discursive representation involved in the pretender's imagining itself seems to have motivational force that can lead to pretense actions. By contrast, in reading a work of fiction, the (potentiated) discursive representation involved in the reader's imagining itself seems to lack such motivational force and, usually, does not lead to any such action. I think that the appearance of conflict disappears once we identify the core functional role that remains consistent across these and other imaginative projects. But the fact that imagining seems to exhibit variation not just in its input conditions but in its features and outputs has made philosophers skeptical of whether we can succeed in this task. One of the main insights of the Skillful Action Account is that success lies in thinking of imagining not just as paradigmatically an action but paradigmatically skillful as well. The features and outputs of identifiable types of bodily skillful actions, say, making a certain brushstroke or making a backhand in tennis, vary along with their inputs. If imagining is a type of skillful action then such variation is to be expected rather than marked as indicative of some conflict among its instances (see the rest of this section). Thanks to the anonymous reviewer for pushing me to clarify this point.

We might be able to characterize certain kinds of imagining and we might be able to use such accounts to elucidate imagining's role among some subset of types of imaginative project. But, so this response maintains, it is not possible to provide a characterization of imagining that captures a single unified mental phenomenon, illuminates its nature in an intelligible way, and satisfactorily unifies its instances across sufficiently diverse types of project. Much of the literature on imagination accordingly adopts a piecemeal approach, some explicitly in light of the worries voiced by Walton and Kind (see, for instance Langland-Hassan 2020; *cf.* Wiltsher 2019a,b, 2023).

Luckily, this is not the only response to the problem. One viable response is to treat the problem as gesturing towards core desiderata for accounts of imagining that aim to unify its instances across types of imaginative project. I count at least two core desiderata. First, a unifying account of imagining should be intelligible in the sense that it is useful for theorizing about imagination and provides some insight into the nature of imagining. Call this the "intelligibility desideratum." Arguably, a source of the discouraging response to the Unity-in-Multiplicity problem is a recognition that accounts that aim for unity tend to give up intelligibility and this undercuts their potential explanatory power (*cf.* Dorsch 2012). An adequate characterization of imagining should capture its essential features, specifically what distinguishes it as a mental activity. What is more, an adequate characterization should specify what features of imagining in particular make it well suited to contribute to the completion of imaginative projects. This leads to the second desideratum on unifying accounts, namely, such accounts should provide a characterization of imagining that explicates the contribution it makes to completing imaginative projects. Call this the "explication desideratum." Arguably, another source of the discouraging response to the problem is a recognition that accounts that aim for unity often end up providing a satisfactory explication of imagining's contribution to only a few types of intimately related imaginative projects at the expense of covering types of project that differ significantly from those it treats as central. An adequate characterization of imagining should show how types of imaginative project generally hang together as types of *imaginative* project.

In this section, I argue that the Skillful Action Account of imagining has the resources to satisfy both desiderata and is thus in a position to resolve the Unity-in-Multiplicity problem. To mount my argument, I draw attention to five parameters common to the performance of skillful action through a series of parallels between skillful bodily action and active imagining. A single type of skillful action can tolerate variation in the values of these parameters across its instances, including values that conflict with the values of parameters of another instance of that same act-type. Such variation and conflict across instances are fully consistent with

type-individuating these actions. One need only identify the means-end order specific to the type of action. I suggest that the same applies to active imagining: its instances exhibit variation in the values of these same parameters, including values that conflict with the values of parameters of other instances of active imagining. Accounting for the conflicting functions ascribed to imagining across types of imaginative project by appeal to the parameterization of skillful action satisfies the explication desideratum. At the same time, the Skillful Action Account sets out the act-type common across such variation and conflict: potentiating (/generating an intentional action acorn for) the construction of a discursive representation and selection of its content as a means of completing an imaginative project. This specification of the act-type satisfies the intelligibility desideratum.

Consider a skillful action like playing piano. On the Skillful Action Account of Imagining, playing piano and active imagining are species of a common genus, namely, skillful action. To unpack this claim, consider the following five parameters on skillful action:

(i) variation in properties across performances;
(ii) variation in causes across performances;
(iii) variation in effects across performances;
(iv) practical knowledge of a performance;
(v) propriety of a performance;

Let us consider each in turn. Focus on playing piano. The properties *(i)*, causes *(ii)*, and effects *(iii)* of individual acts of, say, making keystrokes, pressing foot pedals, adjusting dynamics and tempo, reading sheet music, and so on vary depending at least in part on the end for which these acts are means. Focusing on keystrokes, a keystroke can be slow, quick, careful, careless, and so on. Beyond adverbial properties, keystrokes can involve one set of bodily movements on one occasion and a contrary set of movements on another, be expressive of a particular style on one occasion and expressive of a contrary style on another, and so on *(i)*. Making a keystroke can result from deliberating or from responding to an appropriate cue, say, hearing the preceding note or glimpsing the sheet music *(ii)*. And making such a stroke can result in or come to (at least partly) constitute, say, performing a *glissando,* playing a certain chord, conveying the mood of a piece, and so on *(iii)*. In the bodily case, then, a single type of skillful action like making a keystroke as part of playing piano can exhibit significant variation across its instances at least in part due to those instances constituting means that are less generic and more determinate than the determinable "making a keystroke" to ends that are less generic and

more determinate than the determinable "playing." These more specific and determinate means (making *this* keystroke) and ends (performing *this glissando*) are subsumed under more generic and determinable means (making *a* keystroke) and ends (playing) by virtue of a corresponding subsumption of specific practical knowledge of how to perform a *glissando* on this occasion under the more generic practical knowledge of how to play.

In addition to this outward variation, then, there is significant variation in the details of one's practical knowledge *(iv)* of the acts involved in particular instances of playing piano. As a result, there is significant variation in the propriety of those acts *(v)*. Whether an individual keystroke results in or is properly also the performance of a *glissando*, the playing of a certain chord, a conveying of a certain mood, and so on depends on the agent's knowing how to perform these wider actions as part of the exercise of her practical knowledge on the relevant occasion. The specific practical knowledge that the pianist brings to bear on making a certain kind of keystroke for the sake of performing a *glissando* differs from the specific practical knowledge she brings to bear on making a certain kind of keystroke for the sake of playing a certain chord differs from... Each making of the keystroke differs in light of the pianist's exercising different bits of practical knowledge of the end to which that stroke serves as a means. But even practical knowledge of just the keystroke in isolation will differ across instances, since such knowledge incorporates appropriate respecification of the pianist's intention in light of how her performance is currently unfolding (Small, 2012: 158) *(iv)*. Importantly, the pianist's acquiring actionable knowledge-how to make the relevant keystrokes depends at least in part on sufficient practice performing such strokes in accordance with the principles internal to the domain of her skill, say, playing classical music. Only once she has acquired this knowledge-how can she put it to use in exercising her practical knowledge.[88] Such knowledge-how is chiefly what establishes what is appropriate for any given performance by virtue of reflecting the standards of the musical artform *(v)*.

Recall that the agent's practical knowledge comes to constitute her performance by way of constituting the means-end order which that performance embodies. In the cases under consideration, the most general practical knowledge she exercises in any given performance is that of playing by means of making keystrokes. Such knowledge considered as standing actionable knowledge-how comprises a significant portion of her skill. That is, such standing knowledge forms a significant portion of the capacity itself. Her exercise of more specific

[88] Before then, when a pianist is in the process of acquiring the relevant knowledge-how, her performance is appropriate relative either to her intention to learn how to play and her practical knowledge thereof or to a teacher's intention to teach and his practical knowledge of guiding her to perform thusly (Small 2014).

practical knowledge, in turn, constitutes her making particular keystrokes that result in or come to constitute, say, her performance of a *glissando* as part of an instance of painting. That more specific practical knowledge considered as standing actionable knowledge-how forms more determinate parts of a portion of her skill that I have been considering. That is, such standing knowledge forms the increasingly determinate structure of the portion of the capacity under consideration. Recalling chapters 1 and 2, such standing knowledge forms part of the particular shape *in concreto* of the pianist's skill for playing. It is thus the skill itself and its exemplification in action as a function of the exercise of the relevant bits of both generic and specific practical knowledge, then, that possess these parameters and their values, respectively.

Variation across performances of an action—including performances exhibiting values of parameters *(i)-(v)* contrary to each other—in no way prevents identifying those performances as tokens of the same type. Type-individuating the relevant action depends on identifying its role as a means across the range of wider actions in which its instances are embedded. Although I have so far been focused entirely on a single aspect of playing piano, namely, making keystrokes, what I have said thus far applies to the other acts constitutive of the musical artform. Moreover, recalling chapter 1, I take it that the paradigmatic expression of the skill is actually playing. The means-end order specific to playing piano, then, is the agent's at least minimally aesthetic expression through sound generated by keystrokes and the like as a means of participating in the relevant musical artform.[89] Her capacity to play is constituted at the highest level of generality by her actionable knowledge-how to make such expressions. The musical artform in which some instance of play counts as participating specifies the principles that, when incorporated into the pianist's repertoire of actionable knowledge-how, contribute to establishing what is appropriate for that bit of play. Importantly, I take the

[89] One might worry that this precludes infants and those uninterested in participating in any artform from being able to paint, since, arguably, in generating sounds through keystrokes, neither have the end of participating in such an artform. In response, the former involves the exercise of the parent's or teacher's practical knowledge of playing piano in guiding the infant rather than any practical knowledge of playing on the part of the infant. The infant has no such knowledge. The latter is a limiting case. If the uninterested person's act of generating sound through keystrokes is genuinely an instance of playing piano then it is a case of playing purely for entertainment or for its own sake. And such cases are derivative of participation in the artform. Consider playing tennis for fun or just because you feel like it. For such an act to count as a genuine instance of playing the game, the game and the social praxis around it have to exist. This is true even if you yourself have no awareness of the game or the social praxis, say, because whoever taught you never made it apparent that there was such a game or praxis. It is possible to fashion a racket and ball and make movements identical to those of a tennis player's without having any awareness of the game. But, I maintain, such a case is not one of playing tennis. Rather, this case is one of unknowingly reinventing the game.

relevant set of musical artforms to include those setting out the principles for distinct genres.[90] Any action worth counting as an instance of playing piano will involve the generation of sound through keystrokes and the like in a way that is informed by a social praxis that takes such generation as its primary means, provides at least the potential for mastery of such generation, and is at least minimally concerned with aesthetic value.

The points laid out in the previous three paragraphs apply directly to active imagining. Recall our pretender playing 'the floor is lava'. Regarding parameters *(i)-(iii)*, our pretender's imagining can be slow, quick, careful, careless, and so on. Beyond adverbial properties, imaginings differ in their "format"[91] as well as in their attitudinal and phenomenological properties. Our pretender might potentiate the formation of a discursive representation of the lava that activates perceptual recognitional skills and, thereby, has a perceptual or sensory phenomenal character or "format," that is, "mental imagery" (Van Leeuwen, 2013: 222; Nanay, 2010, 2016, 2021, 2023). Her imagery might be thin and sparse or richly detailed and she might integrate it into her perceptual field (Van Leeuwen 2011; Brown 2018). Or our pretender might potentiate the formation of a linguistic or symbolic representation conveying some state of affairs that she takes to be (merely) fictional or non-actual, say, that, within the pretense, she is surrounded by lava. That is, she might engage in "attitudinal imagining" (Van Leeuwen 2013; Langland-Hassan 2020). Or she might potentiate the formation of a discursive representation simply of the lava itself and treat what she then recognizes as fictional or non-actual. That is, she might engage in "objectual imagining" (Yablo 1993) *(i)*.[92] Our pretender's imagining might be the result of deliberation or might be a response to an appropriate cue, say, being told by another player "look out! The floor

[90] Some genres might be such that overt imagining becomes part of the means of their expression, say, jazz. Similarly, composition of musical pieces might take overt imagining as a means (*cf.* footnote 72). For the case under consideration, I am simply assuming that the agent is playing a piece she knows well but did not compose and is not using her play to represent anything that is not already "there" on the sheet music. Under these assumptions, the performance in question might be aptly described as "lacking in imagination" or "unimaginative." I take it that masterful pianists are often overtly imagining in playing pieces that others have composed by expressing through play a nuanced mood that is not already "there" in the sheet music.

[91] I use "format" in quotes because, on the powers framework, only discursive representations that have actually been constructed can differ in their format. *De agendo* representations do not have a format because they do not represent by way of standing in or going proxy for their objects. In cases of covert mental action, the potentiation of the construction and use of differently formatted discursive representations can have different downstream effects. Potentiation of the construction and use of an image or a sentence in speech, for instance, will prime perceptual recognition where such priming has a phenomenal quality to is sufficient for imagery. By contrast, potentiation of the construction and use of a sentence in written language might well prime recognition of its content for further use in the absence of priming perceptual recognition.

[92] I will admit that I am not quite sure how an agent can objectually imagine on the powers framework except by way of potentiating the formation of a fragment of a discursive representation the completion of which would convey some state of affairs as fictional or non-actual or something like a proposition or Fregean "Gedanke" (thought).

is lava!" Similarly, were an agent imagining an event in which he or another is surrounded by lava as part of engaging in a fiction, his imagining might likewise be quick, slow, careful, careless, and so on. Such imagining might likewise be imagistic, attitudinal, or objectual. And such imagining might be the result of deliberation or might be a response to an appropriate cue, say, reading, "the fires below awoke in anger, the red light blazed, and all the cavern was filled with a great glare and heat" (Tolkien, 1993: 1238) *(ii)*. Our pretender's imagining will lead to further imaginings in keeping with the goings-on of the pretense and will potentiate pretense actions motivated by what she imagines. By contrast, our fiction reader's imagining himself or another surrounded by lava as part of engaging with a fiction should not result in the performance of any actions by which he means to affect the goings-on of the world of the fiction, lest he lapse into pretense *(iii)*.

Parameters *(iv)* and *(v)* bring the skillful nature of active imagining to the fore, specifically as a means of principle-driven completion of imaginative projects. Our pretender's practical knowledge *(iv)* that she is imagining that the floor is lava and why, namely, to carry on the pretense, is at a minimum the application of her knowledge-how to attribute non-factual things to herself, to attribute the relational property of being surrounded by lava, and to play 'the floor is lava'. Application of this knowledge-how in her act of potentiating the formation of the relevant discursive representation is such that it motivates and thereby results in her performing pretense actions appropriate to carrying on the pretense. Her acquiring this knowledge-how depends in the first instance on practice imagining across types of project, including practice engaging in pretense. Indeed, the development and subsequent shaping of the pretender's capacity to imagine is a process of learning through guided practice in joint imaginative acts, unguided practice in solo imaginative acts, and performance in the completion of the relevant imaginative project (Small 2014). Focusing on pretense, infants engage in solo pretend play only for short bursts and can play for extended periods only thanks to scaffolding from others (Nielson and Christie 2007; *cf.* Hess 2006). By the time of adolescence, agents are in principle able to engage in pretense indefinitely, as any avid player of Dungeons & Dragons can attest.

The preceding is true of engaging with fiction as well. Recall our fiction reader. The fiction reader's practical knowledge *(iv)* that he is imagining Frodo's and Sam's approach to Mount Doom and why, namely, to engage in Tolkien's *Return of the King*, is at minimum his application of his knowledge-how to attribute several complex properties to non-existent characters, to relate events that never actually occurred to each other in a narrative order, and to read epic high fantasy. Application of this knowledge-how also depends in the first instance on practice imagining across types of project, including practice engaging in fiction. The development and subsequent shaping of the fiction reader's capacity to imagine is a process of learning through guided

practice in joint imaginative acts, unguided practice in solo imaginative acts, and performance in the completion of the relevant imaginative project. In learning to read fiction, the agent learns more than how to parse the relevant sentences and words in context. He learns how to immerse himself in a world that is not his own (usually) *without being motivated to act so as to affect the goings-on of the fictional world*. He learns how to represent episodes in the fiction and relate them to each other in a narrative structure, how to attribute mental states to others (sometimes on the basis of being told what those others are thinking), and so on all with the recognition that the subjects of his mental actions are not, strictly speaking, supposed to exist or be influenceable by his actions. He acquires this actionable knowledge-how to engage in a distinctive mode of thought (or of entertaining imagery) through guided and unguided practice in reading fiction, engaging other imaginative projects, and reading non-fiction.

Relating the capacity to imagine to the propriety of imaginative performances *(v)*, part of what our pretender acquires through practice is grasp of the principles that guide her pretense behavior. Each type of imaginative project brings with it some such principles that partially define the corresponding domain of the skill. These principles govern not only the production and transformation of imaginative states and their contents but the production and transformation of corresponding non-imaginative states and behaviors involved in the successful completion of the relevant imaginative project as well. For example, there are arguably two primary principles at work in pretense (Gendler 2003, 2006a,b). First, actions performed within the pretense mirror the causal and logical structure of the corresponding non-pretense behaviors were the events of the pretense to occur at the actual world. Second, the goings-on of the pretense, including pretense actions involving real-world behavior, are taken to have effects quarantined to the within the pretense. Our pretender's imagining is usually formed and updates in accordance with these mirroring and quarantining principles and it is these principles that she exploits in performing pretense actions, injecting new content into the pretense, and intentionally violating those very principles.[93] She becomes proficient in pretending in part by internalizing both principles through guided and unguided practice, where the outcomes of her attempts therein serve as

[93] Gendler (2003, 2006b) refers to such violations as disparity and contagion, respectively. An agent engages in imaginative disparity as part of pretense when, for instance, she pretends that a toy horse that she is feeding within the pretense is making chewing sounds and subsequently makes those sounds herself. The spatial relations between herself and the horse are mirrored within the pretense but where the sound is coming from within the pretense diverges from where it is coming from in the actual world. An agent engages in imaginative contagion in pretense when, for instance, pretending to be a person of an opposing political party leads her to soften her negative opinion of members of that party. Acting out their way of life or policies can give her a greater appreciation for their point of view, leading her to extract that appreciation out of the pretense and make use of it in her actual dealings with members of the opposing party.

feedback. Such internalization just is her acquiring corresponding bits of fairly generic actionable knowledge-how to pretend. That is, this process quite literally shapes her capacity to imagine.

Once incorporated into her capacity to imagine, the principles guiding the completion of a type of imaginative project come to underwrite the agent's acts of potentiating the construction and use of the appropriate content *(v)* in completing tokens of the relevant type of project. As with other types of skillful action, the proficient imaginer puts her hard-won knowledge-how to use in exercising her practical knowledge. And because exercise of her practical knowledge constitutes her imaginative performance by way of constituting the means-end order which that performance embodies, it is the skill itself and its exemplification in action that possess these parameters and their values, respectively.

Again, the preceding applies to engagement with fiction as well. Walton (1990) identifies two primary principles for generating fictional truths, i.e., facts which hold of the relevant fictional world(s), that a competent audience will apply when engaging with the relevant fiction. First, if propositions $p_1,...,p_n$ are ones whose fictionality a piece generates then another proposition, q, is fictional relative to that piece just in case if $p_1,...,p_n$ were the case then q would be the case. Walton calls this the Reality Principle (RP) (1990: 145). A fictional truth that is generated directly is one that is asserted by way of features of the piece itself given the conventions and, thus, social praxes surrounding the production of such pieces. In our running example, directly generated fictional truths would be those expressed on the page by Tolkien. RP sanctions further indirect generation, where the fictionality of a proposition relative to a piece is inferable from another in accordance with RP even if the latter is not itself generated directly by the piece. For instance, our reader can recognize that it is fictional that Frodo and Sam are sweating by inferring from the fictional truths that hobbits have physiology at least not unlike ours and that the temperatures inside Mount Doom are likely to be similar to those at the mouth of an erupting volcano at the actual world. Neither of these fictional truths are (to my knowledge) generated directly by *Return of the King*.[94] Second, if propositions $p_1,...,p_n$ are ones whose fictionality a piece generated directly then another proposition, r, is fictional relative to some piece just in case it is mutually believed in the author's society that if $p_1,...,p_n$ were the case then r would be the case. Walton calls this the Mutual Belief Principle (MBP). Walton notes that the ways by which fictional truths are generated directly is conventional. He also notes that neither RP nor MBP are meant to be universal principles. They can

[94] Walton notes issues that may arise in cases where contradictory fictional truths are directly generated by a piece (1990: 147-150). I leave these issues aside for the sake of space but wish to follow Walton in ignoring rather than banishing the fictional truths generated by such explosions.

be violated, like the principles guiding pretense. The important point is that our reader's imagining is usually formed and updates in accordance with RP and MBP and it is these principles that he exploits in elaborating the fictional world of Middle Earth, in entertaining himself by wondering whether hobbits braid their feet hair, and so on. He becomes proficient at engaging in fiction in part by internalizing these principles through guided and unguided practice, where the outcomes of his attempts therein can be compared to public expressions of others' and, thus, indirectly serve as feedback. Such internalization, like that involved in pretense, quite literally shapes his capacity to imagine.

Again, once incorporated into his capacity to imagine, the principles guiding the completion of a type of imaginative project come to underwrite the agent's acts of potentiating the construction and use of the appropriate content *(v)* in completing tokens of the relevant type of project. And, again, because exercise of his practical knowledge constitutes his imaginative performances by way of constituting the means-end order which that performance embodies, it is the skill itself and its exemplification in action that possess these parameters and their values, respectively. The differences between our pretender and our fiction reader, specifically the differences between the motivating features of the former's imagining and the lack thereof of the latter's, stem from differences in the principles appropriate to the respective domains of each type of project. In particular, the mirror and quarantining principles in the domain of pretense and the lack of such principles in the domain of engaging with fiction facilitate imaginings with distinct functional profiles in proficient imagining agents exercising their capacity in either domain. Acts of imagining backed by these principles consist in potentiations of the formation of discursive representations that have positive or negative valences or the character of directives or further *de agendo* representations for appropriate pretense actions. By contrast, acts of imagining backed by RP and MBP result in the potentiation of the formation of discursive representations that do not have these valences or characters. One possible exception is interactive fiction, which is arguably a kind of pretense (*cf.* Walton 1990). Recalling chapters 1 and 2, because the Skillful Action Account has it that such principles come to shape the capacity to imagine in the form of embodied, actionable knowledge-how, this means that the capacity itself is structured as a complex of dispositions for exercising such knowledge-how in appropriate contexts and in light of the imagining agent's intention to imagine thusly. It is the structure of the complex—fashioned through inculcation into the relevant domains—and activation of the appropriate parts of it *qua* exercise of the imagining agent's practical knowledge that allows individual exercises of the capacity to imagine to have seemingly inconsistent properties when we consider the capacity independently of any of its exercises. The insight of the Skillful Action Account is that the same is true of other capacities worthy of being

called skills. As the case of playing piano shows, any two manifestations of a capacity for skillful action might conflict along their more determinate dimensions because such capacities comprise distinct tracks consisting of more determinate capacities for actions exemplary of the skill and because each manifestation might utilize distinct tracks.

I am now finally in a position to show that the Skillful Action Account can satisfy the intelligibility and explication desiderata on unifying accounts of imagining. Starting with the explication desideratum, recall that an adequate unifying account should provide a characterization of imagining that explicates the contributions it makes to the completion of imaginative projects. The parameters of skillful action show how the account can do this. First, if, as the account has it, active imagining is a skillful type of action then variation in the properties, causes, and effects of its instances is to be expected. This includes cases in which the values of the parameters in one performance conflict with the values of the parameters in another performance. Second, the agent learns through practice the principles specific to each type of imaginative project. She then applies what she has learned in exercising her practical knowledge in imagining. In possessing this knowledge, she potentiates the construction and use of discursive representations that are systematically attuned to the corresponding imaginative project that she intends to complete. Conflict in the role of playing piano across genres or styles is resolved once it is acknowledged that the pianist is in one context intending to do one thing in playing and is in another context intending to do something else in playing. What is shared across these contexts is her intentionally performing the same generic act as a means to these distinct ends under the self-imposed guidance of the relevant principles. By the same token, both conflict in the role imagining plays across types of imaginative project as well as conflict in the features that are supposed to allow it to play its roles are resolved once it is acknowledged that the sufficiently proficient imaginer is in one context intending to do one thing in imagining, say, *p* and is in another context intending to do something else in imagining, say, *q*. What is shared across these contexts is her potentiating(/generating an intentional action acorn for) the construction of a discursive representation and her selection of its content as a means to completing distinct imaginative projects under the self-imposed guidance of the relevant principles.

Moving on to the intelligibility desideratum, recall that an adequate unifying account should be intelligible in the sense that it is useful for theorizing about imagining and provides some insight into its nature. By laying out the means-end order specific to imagining as a mental action and by fleshing its characterization out with the five parameters of skillful action, the Skillful Action Account can do this. The claim that active imagining is paradigmatically a type of skillful mental action is straightforwardly intelligible. The claim that this

mental action consists in being primed to form a discursive representation, to select its content, and to make use of that representation as a means to pretending, engaging with fiction, mindreading, and so on is illuminating of the nature of imagining. The addition of the five parameters enriches this characterization. In particular, the skillful dimensions of imagining come into focus by considering how acquiring knowledge-how to engage appropriately in different types of imaginative project gives determinate shape both to the capacity to imagine and to essential features of its individual exemplifications. It becomes clear how, through practice, the agent's practical knowledge of those imaginings becomes increasingly nuanced and sensitive to the demands of the corresponding imaginative project. Finally, the characterization of imagining provided by the Skillful Action Account is ripe for theorizing. And this brings us back to the Unity-in-Multiplicity problem. I have so far provided an extremely sketchy programmatic roadmap for addressing the problem. Fully addressing it requires, first, laying out the principles specific to each type of imaginative project and, second, showing that sufficient practice constructing discursive representations in accordance with those principles is in fact how agents become proficient pretenders, fiction readers, mind-readers, and so on. Such is the plan for an empirically informed philosophical research program suggested by the Skillful Action Account of Imagining.

3.4 Applying the Account Part II: The Intentionality of Imagining

So far, I have been concerned primarily with the nature of imagining as a kind of mental process, that is, as a type of a mental action. However, because the activity that imagining consists in is one of potentiating(/generating an intentional action acorn for) the formation of discursive representations, a complete account of its nature requires discussion of its intentionality. As was noted in Section 3.2, the Skillful Action Account holds that the agent selects the content of what she is primed to represent in imagining in light of the end she has in potentiating the formation of that representation. In this section, I show that the account adopts a qualified form of what Munro and Strohminger (2021) identify as "intentionalism." According to intentionalism about the content of imagination, "whenever you intend to imagine something and act on that intention, the content of your intention about what to imagine determines the content [that] you imagine" (11848).[95] This formulation of intentionalism is unrestricted. By contrast, on the Skillful Action Account, there are limits to what an agent can imagine that stem from the limited nature of human practical cognition in general. Just as successful action requires sufficient knowledge-how, sufficient sensitivity to the context of

[95] Intentionalism about the content of imagination can be traced at least as far back as Wittgenstein (1980) and is adopted by Fodor (1975), Noordhof (2002, 2018), McGinn (2004), Dorsch (2012), Langland-Hassan (2016), Balcerak Jackson (2018), and Kind (2019).

action, and efficacious performance, so too does successfully potentiating the formation of a discursive representation of what one intends in imagination. On the qualified form of intentionalism put forward by the Skillful Action Account, only when the agent's practical cognition of her act of imagining achieves the status of an exercise of practical knowledge does she imagine what she intends. This account of the intentionality of imagining further enriches the characterization of the nature of active imagining proposed by the account.[96]

The relation between intention and content in active imagining is, presumably, an aspect of the freedom of imagination. Some form of intentionalism is thus the norm in the literature on imagination. However, as Munro and Strohminger point out, there are cases where it seems that what is imagined diverges from what the agent intends to imagine. Let us focus on one such case of imagistic imagining inspired by Wittgenstein (1980). A tourist visits both King's College and Trinity College. Later, she mixes up her past experiences in attempting to episodically remember the visits. She now misremembers her tour of Trinity College as being her tour of King's College. She then intentionally imagistically imagines what she takes to be King's College on fire. Intuitively, the agent misimagines King's College on fire. Her image mischaracterizes King's College as Trinity College and, thus, fails to appropriately represent the former. I grant that this case counts as a genuine instance of misimagining. It thus gestures at a genuine limit on the freedom of imagination that is tied to the etiology of the imagery involved.[97]

The question, then, is how to square the possibility of misimagining with the intentionalist dimension of the freedom of imagination. According to the Skillful Action Account, at its core, what goes wrong in cases of misimagining is a failure of an attempted imagining to match the agent's practical cognition of what she is doing. Such failures can occur in more than one way (see Appendix C). The agent can i) fail to know how to imagine (appropriately), ii) fail to know some fact upon which her success in this instance depends, or iii) make so-called "mistakes in performance." In the good case, the agent's fixing the content of an imagining occurs as

[96] This is not all there is to an account of the intentionality of imagining. In particular, the Skillful Action Account is neutral regarding whether there is a semantics for imagining or what such a semantics might look like. There are several accounts of the (formal) semantics of imagining on offer. My preferred account can be found in Berto (2017, 2018, 2023) (see also footnote 81). According to Berto, what is represented in imagining are sets of possible or impossible worlds that are accessed by virtue of the agent's act of imagining and where what is posited in an act of imagining is true (see also Giordani, 2019; *cf.* Casas-Roma, Rodríguez, and Huertas, 2019). I take it that imagining's intentionality makes it exploratory of temporal and modal space.

[97] Another case that Munro and Strohminger (2021) consider involves only touring Trinity College under the false belief that one is touring King's College and then attempting to imagine King's College on fire. This example strikes the authors as another case of misimagining. However, one might think that their intuition about the case presupposes wide mental content. An intentionalist proponent of narrow content might deny that this case is one in which the agent misimagines. To remain neutral on the narrow-wide mental content debate, I decline to use the example.

part of the exercise of her practical knowledge in intending to complete the relevant imaginative project. However, like all cognition, practical cognition is fallible such that it sometimes fails to reach the status of knowledge. An agent, in acting, can think they are doing one thing when in fact they are not.

The reasons for such failures in the bodily case mirror the ways that misimagining can occur. First, an agent can think he knows how to do something that he in fact does not know how to do, and, upon attempting to exercise what he takes to be a bit of actionable knowledge-how, fails to do what he intends. For instance, he might think he knows how to tie a full Windsor knot but, in fact, does not know to make a second loop on the left side of the tie. His attempt at exercising what he takes to be knowledge-how to tie a full Windsor will at best result in a half Windsor. Second, an agent can know how to do something and set out to do it but fail to know some fact that the success of his performance depends on. For instance, he might know how to operate a manual water pump but not know that there is a hole in the water suction line. His attempt to pump water will likely result in pushing a lot of air out of the pump.[98] Finally, an agent can know how to do something and can know everything that he needs to in order to succeed in performing the relevant action on some occasion but, nonetheless, his performance fails to realize his intention. For instance, he might know how to call an elevator and know that button *A* calls the one that he means to call. Nonetheless, in reaching for button *A* he might miss and press button *B*. In all three cases, what ends up happening fails to match the agent's practical cognition.

Returning to imagining, for each way in which what ends up happening fails to match the agent's practical cognition, the agent misimagines. In the King's College case, the tourist's failure falls within the second category. The tourist knows how to imagine buildings on fire but does not know that the imagery of burning buildings that she has are those belonging to Trinity College rather than those belonging to King's College. In particular, she does not know that she has misattributed her memory-based imagery of Trinity College to her past experience of touring King's College. This fact turns out to be pertinent to the success of her attempt to imagine King's College on fire. She can only correct for the mistake if she is made aware of it. Now, consider an agent, an idle imaginer, who knows how to imagine that the floor is lava but does not know that she has been invited to play a game of 'the floor is lava'. Instead, she thinks that she has been invited simply to entertain imagery or the thought of the floor's being lava and it is her taking herself to be so invited that prompts her to entertain the relevant imagery or thought. In this case, the idle imaginer fails to know that the end she has been

[98] Note that failures of knowledge-how and of knowledge of facts required for successful performance can be as generic or specific as you like. An agent can fail to know how to tie a tie, how to tie a full Windsor, or how to tie *this* tie. Similarly, an agent can fail to know that manual pumps exist or what a pump even is, that pumps function to move fluids, or that *this* pump is broken.

invited to take up is engaging in the pretense. She is thus liable to misimagine because she is likely to fail to update her imagining in response to her or others' pretense actions. She will likely fail to update her imagining in response to another player tossing a throw pillow onto the floor because she will likely fail to recognize that action as that player's making a move in the game. The idle imaginer's liability to fail to update in this case suggests that failures to know facts required for successful imaginative performance concern not just content or features of the context of action but also the point of imagining in a given instance.[99]

There are also failures of knowledge-how that result in misimagining. Such failures take at least two forms: an agent's failing to know how to imagine something and her failing to know how to imagine such that she completes the corresponding imaginative project. Regarding the first form, a greenhorn orchidologist might not know how to imagine a king-in-his-carriage orchid because she is completely unaware of the species. Upon being told about the orchid, she might try to imagine it and form imagery of an especially glorious-looking flower belonging to the *Cypripedioideae* genus. But she would be misimagining the orchid in this case, since king-in-his-carriage is of a genus of orchids with a peculiar look (Figure 3.1).[100] Regarding the second form of failing to know how to imagine, a novice player of a game of pretense might not know how to imagine appropriately for engaging in that game. She might be just starting to learn the rules of 'the floor is lava' and, upon seeing another player toss a throw pillow onto the floor, not know to update her imagining. She might then attempt to avoid touching the throw pillow. She might think it is, say, a lava-resistant naval mine. In such a

[99] One might worry that this is not a genuine case of misimagining. Rather, the idle imaginer succeeds in imagining what she intends to imagine, namely, the floor's being lava. What goes wrong is that she fails to have the appropriate end, namely, to engage in the pretense. In response, the failure to have the appropriate end redounds on the propriety of the properties of the discursive representation that she is primed to construct. On the Skillful Action Account, in failing to have the appropriate end, the idle imaginer is primed to construct a discursive representation that is not appropriately valenced or fails to have the character of a directive or *de agendo* representation for the relevant pretense actions. It is thus a case of misimagining so long as the agent is indeed idle. The potentiated representation might have these properties while it still being the case that she misheard her would-be playmate were she to adopt the intention to play a game of 'the floor is lava' independently of the invitation. In such a case, she would not be misimagining. Yet, such a case strikes me as analogous to a Gettier-case. Thanks to the anonymous reviewer for pushing me to clarify this point.

[100] This does not mean that an agent who knows what king-in-his-carriage orchids look like is beholden to imagining them with their peculiar look. Having the correct imagery does not depend on resemblance with the thing imagined. After all, an agent who knows what king-in-his-carriage orchids look like can intentionally imagine them as though they were a species of the *Cypripedioideae* genus. Similarly, an agent who knows what King's College looks like can intentionally imagine King's College as though it was Trinity College and on fire. Typically, what is going on in misimagining is that the agent is mistaken about how to imagine or about what they aim to imagine, and their mistake obstructs their ability to represent in accordance with their intention. An agent who is not mistaken about her lack of knowledge about what King's College looks like can imagine it on fire and it is important to her ability to do so that, in doing so, she knows that she is not necessarily trying to achieve resemblance in her imagery.

case, she would be misimagining the state of the game: unlike practiced players, she would not select an appropriate content, e.g., that an outcropping has just emerged from the lava flow.

Figure 3.1 King-in-his-Carriage Orchid
A king-in-his-carriage orchid (*Drakaea glyptodo*). The *Drakaea* genus of orchid, commonly known as "hammer orchids" share this peculiar shape with an insectoid labellum to trick species of thynnid wasps into pollinating them. Photo credit: Mark Brundett (2010).

Finally, there are mistakes in performance in imagining. An agent can know how to imagine in both of the relevant senses and can know everything that she needs to in order to succeed in her act of imagining on some occasion but, nonetheless, her performance for some reason fails to match her practical cognition. Suppose our pretender is a practiced player of 'the floor is lava' and knows that another player has just tossed a throw pillow onto the floor. She might nonetheless fail to update and try to traverse the space without making use of the pillow. In such a case, she'd be misimagining the state of the game, failing to imagine that an outcropping has just emerged from the lava flow. She might well feel silly for making this mistake. After all, she knows how to play and would otherwise immediately update her imagining upon seeing the throw pillow get tossed. Still, the content of her imagining failed to match her practical cognition with respect to carrying on the pretense and this mistake might well cost her her life in the game.

The three kinds of failure just listed each constitute a way that the agent misimagines. And each way of misimagining can occur via distinct routes. Each results in the agent's failing to be primed to represent what she intends to. In cases of failing to know how and cases of failing to know some fact upon which her success depends, the agent's expression of her intention would "fall to the ground." Such falling to the ground is not an intention's forever going unsatisfied or its being contradicted but, rather, its currently not being executed by

whatever means the agent is taking (Small, 2012: 143-150). Other things being equal, the agent can pick her statement back up if 1) it is not too late to correct her imaginative act or it is not a one-off, 2) she does not change her mind, and 3) she is made aware of the relevant failure.[101] Indeed, upon learning that tying a full Windsor involves making a second loop, the agent tying his tie will undo the half Windsor and try again. Similarly, upon learning that there is a hole in the water suction line, the agent operating the pump will try to fix the hole or go find some other source of water (Small 2012: 147). In mistakes in performance, what an agent would express in saying what she is doing is "directly falsified." Such direct falsification is likewise not an intention's forever going unsatisfied or its being contracted. Rather, it is the failure of the agent's performance to match her practical cognition. And she can rectify such a failure: upon learning that he had pressed the wrong button, the agent calling the elevator will press button *A*. In all of these cases, the agents go on to satisfy their intentions by adopting the appropriate means once they have learned that what they actually did was not appropriate.

Holding the three conditions on picking up an intention fixed, let us reconsider what happens in the cases of misimagining so far discussed. In the King's College case, upon learning of her mistake, the tourist will use an image search engine for "King's College" to relearn what it looks like, will attempt to correct her memory on her own, will (re)specify her intention to be that of imagining King's College Trinity-College-wise and on fire, etc.[102] Upon learning that she is been invited to play a game of 'the floor is lava', the idle imaginer will shift

[101] The phenomenon of intentions falling to the ground and of the agent's being able to pick them back up are both intimately connected to the so-called "broadness" and "openness" of the progressive (Falvey 2000; Galton 2006; Thompson 2008; Small 2012). Unpacking the broadness and openness of the progressive goes beyond the scope of this chapter. However, see Appendix C.

[102] One might worry about the extent to which imagery must be accurate in order to succeed as an instance of imagistically imagining such-and-such. What if the agent in this case is imagining college-y buildings on fire but nothing particularly distinctive of either King's College or Trinity College? Would the agent count as misimagining? In response, it depends. If the agent intends to imagine King's College on fire then she has to know how to imagine the college on fire in a way appropriate to her intention and she has to exercise that know-how in accordance with her intention. If her imaginative project is one of entertaining herself then the imagery may not need to "contain" anything distinctive of King's College, so long as she stipulates that the college-y buildings on fire are those belonging to the college. If her imaginative project is one of figuring out what the best escape route would be were King's College on fire then the imagery should "contain" enough details distinctive of the college to allow her to infer that such-and-such a route is best. If her imaginative project is one of engaging in pretense with others who know the college well and wish to play in as realistic a pretense as possible then the imagery should "contain" enough details distinctive of the college to allow her to make legitimate moves in the game. And so on. That the appropriate level of accuracy is contextual and based on what imaginative project the agent is engaged in is a feature of the skillful action view that I consider one of its strengths. I also take it to be in keeping with Kind's (2016a, 2016b, 2018), according to which mistaken imagery can nonetheless result in successful imaginings. I agree. What the Skillful Action Account is committed to denying is that an agent can form imagery that fails to close the distance between her and the completion of her imaginative project and still count as having completed that project. Whether the imagery

from merely entertaining the relevant imagery or thought to imagining appropriate for engaging in the pretense. Upon learning what a king-in-his-carriage orchid looks like, the greenhorn orchidologist will either correct her imagery or will (re)specify her intention to be that of imagining a king-in-his-carriage orchid *Cypripedioideae*-wise. Upon learning the relevant updating rules for 'the floor is lava', the novice player will go on to apply those rules to her imagining in continuing to try to play. Finally, upon learning that she had failed to update, the seasoned player will automatically update to include the new outcropping as among the walkable surfaces within the pretense.

In each case where the agent addresses some *prima facie* failure of practical cognition or corrects her performance, her correcting her mistake means that she does not ultimately misimagine. That is, in these cases, the agents successfully imagine King's College on fire, that the floor is lava, and the king-in-his-carriage orchid, respectively. An agent's being in a position to address some *prima facie* failure of practical cognition or a mistake in performance and carry her action to completion suggests another way in which misimagining constitutes a failure to fully exercise the capacity to imagine. Indeed, misimaginings are in general relegated to defective portions of ultimately successful imaginings, instances where the agent overestimates her (current) imaginative ability,[103] instances where the agent gives up before successfully correcting her imaginative behavior, and instances where there is no chance for correction.

By providing a more robust account of imagining's intentionality, the qualifications that the Skillful Action Account makes to intentionalism further enrich the account's characterization of the nature of imagining as paradigmatically a type of skillful mental action. What an agent imagines is what she intends to imagine just in case, other things being equal, her act of imagining matches her practical cognition of that act. Adding this qualified form of intentionalism to the Skillful Action Account of Imagining has the following two results. First, what an agent intends to imagine and, thus, what she ends up imagining are both shaped by the intention with which she imagines, namely, to pretend, engage with fiction, mindread, imagine for its own sake, and so on. Thus, the agent's freedom to imagine is tempered by whatever end for the sake of which she is

---

one forms fails in the relevant respect will depend at least on the project and the context of the agent's attempt to complete it. I likewise take this contextual aspect of the Skillful Action Account to be in keeping with Recanati's (2007, 2010) contexts of evaluation. Thanks to the anonymous reviewer for allowing me the opportunity to elaborate it, if only briefly.

[103] Cases of overestimation include those stemming from lack of the relevant knowledge-how or lack of knowledge of some fact that the success of the relevant imagining depends on. Moreover, should it turn out that we cannot imagine the impossible then such cases also include instances where an agent takes himself to imagine something that it is not possible to imagine, for instance, what a Euclidean triangle whose internal angles summing to 270° looks like. I want to thank Zachary Gabor for this example.

exercising that freedom. Second, an especially important enabling condition for the agent's imagining what she intends is her having had sufficient practice imagining and doing so in accordance with the principles guiding the relevant imaginative project. Just having the needed "ideas" is not enough. The agent must be able to selectively combine elements of those ideas in potentiating the construction and use of a discursive representation and must be able to do so with a view to completing the corresponding project. So, according to the Skillful Action Account, while it is true that in the good case the agent imagines what she intends, what she can efficaciously imagine and what she can genuinely intend to imagine are both constrained by the shape of her capacity to imagine and by whatever imaginative project she is engaged in.[104]

3.5 Addressing Non-Active Imagining and Mind-Wandering

Because the Skillful Action Account treats imagining as paradigmatically a type of skillful mental action, one might expect it to struggle accounting for passive imaginings like automatic imagining, involuntary imagining, or mind-wandering. Neither automatic nor involuntary imagining appear to be types of imaginative project nor appear to occur within such projects. And mind-wandering, at least when understood as "unguided attention" to things other than what is present (Irving 2016), appears to be a type of imaginative project that is by definition undirected by the agent. Indeed, all three appear to happen independently, if not entirely against, the

[104] One might worry about the possibility of misimagining in cases of non–imagistic imagining. So far, the cases I have considered have either been those in which it is explicit that the agent is attempting to imagistically imagine something or those in which it is ambiguous whether or not they are attempting to imagistically imagine. One might thus wonder whether specification of the project constitutes the imagining in cases where the agent is attitudinally or objectually imagining and, if so, how this allows for misimagining as I have described it. After all, there seems to be no distance between specification of the project and its completion in such cases. In which case, there is no distance to close by the agent's taking the relevant means. As it stands, then, the Skillful Action Account might only account for imagistic imagining or might need to provide a distinct account of misimagining in cases of attitudinal or objectual imagining.

In response, I do not take it that an agent's specification of an imaginative project is sufficient to complete it in these cases. Rather, the relevant suitable happening must occur (see footnote 81). That is, the agent must potentiate the formation of a corresponding discursive representation that she treats as fictional or non-actual and which constitutes her locating the relevant set of possible worlds or must represent the relevant object and treat it as fictional or non-actual, where her doing so is partly a matter of her locating the relevant set of possible worlds. Such happenings are distinct from her specifying the project she aims to engage in, since no suitable happening need occur. And, so, there is a distance between specifying the project and completing it such that misimagining can occur, namely, in instances where the agent fails to locate the relevant set of possible worlds. Such misimagining might fail to have any upshots downstream of the mistake if, say, the agent is engaged in reading *Return of the King* and misimagines, say, that Mount Doom is to the north of Barad-dûr. In such a case, the agent will have mislocated the relevant set of possible worlds as that in which Mount Doom is to the north of Barad-dûr. The correct set is that in which Mount Doom is to the southwest of Barad-dûr. Yet, if she never looks at a map of Mordor or is never corrected by anyone, this misimagining might be a contained mistake. If she does look or is corrected, she will know that she misimagined. And she can misimagine, realize this, and correct her imagining without any imagery. Thanks to an anonymous reviewer for pushing me to clarify this point.

agent's will. One might think, then, that to the extent that the Skillful Action Account emphasizes the centrality of the skillful mental action of imagining, it fails to countenance automatic imagining, involuntary imagining, and mind-wandering.

I give five responses on behalf of the account. First, starting with automatic imagining, many such instances might well be responses to an appropriate cue. At least in these instances, automatic imagining is not genuinely passive. This is unsurprising if we take imagining to be skillful. Indeed, skillful action in general depends heavily on automatic action: if the agent had to slow down and think through every movement involved in, say, playing piano every time she sat down to play, she would not be able to master the skill (Wu 2016). Recall our pianist from Section 3.3. Her making a certain keystroke might well be an automatic response to hearing the preceding note or glimpsing the sheet music. Likewise, our pretender playing 'the floor is lava' might imagine the emergence of an outcropping automatically in response to seeing another player tossing a throw pillow onto the floor. Much active imagining is (as it must be) automatic yet skillful, occurring without prior deliberation as a function of habit and in response to the appropriate cue.

Second, moving onto involuntary imagining, some instances of involuntary imagining will turn out to be disruptions in the functioning of mechanisms involved in potentiating construction or in the activation of recognitional skills. In that case, they are not instances of imagining at all—they are defective. Hallucination after a head injury is a case in point. It would at the very least be wildly misleading to describe someone hallucinating after a head injury as imagining. He is at most suffering unbidden imagery as a result of damage to mechanisms involved in potentiating construction or in the activation of perceptual recognitional skills in a way that is difficult or impossible for him to distinguish from genuine perception. It would be more apt to say that he is "seeing things" or "hearing things" with the inflection that the imagery is pathological in part by virtue of its etiology and in part by serving to mislead him. Third, other instances of involuntary imagining will turn out to be non-malfunctional instances of unbidden imagery or thought. Many biological mechanisms are liable to misfire despite not malfunctioning. Indeed, such misfires can be adaptive, say, when a prey detector's being subject to many false alarms brings with it the occasional and much needed meal (Godfrey-Smith, 1992: 297-308). Similarly, the mechanisms involved in potentiating construction or in the activation of recognitional skills are not foolproof. Even if infrequent, they can be activated in the absence of any appropriate stimulus or need. In that case, they will produce unbidden imagery or thought. This might be what happens, for instance, when

people experience unbidden imagery as they are falling asleep, when they dream,[105] or when they suffer an earworm. Such cases are not unlike "mindless" scratching in response to a phantom itch. Such scratching does not remove an actual irritant from the skin but might nonetheless be adaptive if it is reflective of a routine behavior that is likely to activate in response to an actual irritant causing an itch. I also take it that cases of imagery-involving illusions induced, say, in the lab involve (deliberately) causing such mechanisms to non-malfuctionally misfire.

Fourth, still other instances of involuntary imagining will turn out to be action slips. Action slips occur when an action is tokened in response to what would otherwise be an appropriate cue but where performance of that action is inappropriate for the context (Amaya, 2013, 2020; Mylopoulos, 2022). Habitual actions are prone to slips. For instance, an agent can find herself driving her normal route home despite having formed an intention to stop by the grocery store prior to getting into the car. Getting into the car places her in surroundings likely to trigger the habit of taking her normal route. If, as the Skillful Action Account has it, active imagining is the central imaginative kind and is a type of skillful action then we should expect imagining to admit of occasional slips. Some instances of automatic imagining fall into this category, as will some instances of unbidden imagery or thought. In cases of imaginative action slips, the capacity to imagine is likely triggered by a cue for engaging in some imaginative project in the absence of the agent's being engaged in that project. Just like in the bodily case, the agent can become aware of these slips in imagining and can usually correct her imaginative behavior.

Fifth, and finally, the first and fourth responses given to address automatic and involuntary imagining apply as well to mind-wandering. Mind-wandering deserves a more thorough treatment than I can give it here. However, I can provide some reason for thinking that the Skillful Action Account of Imagining can incorporate mind-wandering as an imaginative project. I consider automatic mind-wandering and mind-wandering that starts as an action slip in order. First, recall that, following Irving (2016), mind-wandering is "unguided attention" where such guidance is consciously experienced. Attention is unguided if its movement from one topic to another does not cause the agent to feel "pulled back" from what she was previously focused

[105] It is unclear whether dreaming is to be thought of as a kind of imagining. On the Skillful Action Account of Imagining, if dreaming is not non-malfunctional misfiring of the sub-personal mechanisms involved in potentiating construction or in the activation of recognitional skills then it is likely a distinct mental kind that nonetheless draws on those same mechanisms. Human perception likewise draws on these mechanisms. Because the account presents an action-first non-reductive ontology of imagining, this overlap in the relevant neural correlates or neurocognitive architecture is not sufficient to identify sameness in mental kind. Perception is not a kind of imagining. Arguably, neither is dreaming.

on. Indeed, mind-wandering often lapses into thought about future tasks and wider goals that are more pressing or feel more rewarding to think about (Shepherd 2018). In which case, it is fully consistent with the lack of the feeling of being pulled back that mind-wandering is primarily a matter of automatic habitual or skillful seeking out of behavior that is more efficient or rewarding than what the agent is currently doing. The Skillful Action Account can therefore classify such seeking-out as a subset of a type of imaginative project, namely, (episodic) hypothetical or counterfactual thought about contingent matters of fact (De Brigard, 2014a).

For example, while making dinner, I am also writing this chapter. My thoughts are liable to drift towards the metaphysical structure of action and its relation to imagination as I slice, mince, sauté, and so on. Yet, I do not experience being pulled back from my cooking and I recognize that the task of writing this chapter is, relative to my larger goals, more pressing than eating the meal that I am making right now.[106] In this case, my mind-wandering is my seeking out a more pressing goal automatically in response to an appropriate cue, say, the stress or anxiety I am feeling about completing it. Such cues can be external as well: a bored elementary school student might automatically mind-wander to thoughts of play in response to seeing the school's lawn outside the window. What is important is that the imagining that the student and I engage in is an automatic potentiation of the construction (and use) of a discursive representation whose content is given by the more pressing (or fun) goal we would rather be engaging in, e.g., work or play respectively. That we do not experience such thought as a distraction to be corrected and, so, is attentionally unguided is no bar on our mind-wandering being something we each do, albeit without having to use attentional resources in an attempt to so wander. The paradoxical air around mind-wandering is that it is an action that we cannot perform deliberately. But so too are certain kinds of meditation or mindfulness practices. In all three cases, the key to dispelling the air of paradox is seeing that a lack of guidance or deliberate attempt is consistent with automatic performance.

Similarly, it is consistent with the Skillful Action Account that some instances of mind-wandering begin as action slips. An agent's feeling hungry at work might cause her to think about dinner, where this thinking practically conflicts with her operative intention to stay focused during a work meeting. The agent can then choose to continue engaging in this imaginative project, taking an active role in its development. She might find fantasizing about her upcoming meal a more worthwhile endeavor. Like other slips, one often notices or finds themselves mind-wandering. In such cases, one then either corrects the behavior by refocusing attention to

[106] This might be indicative of some pathological ordering of my priorities. But I suspect it is not an uncommon experience.

the task at hand or elaborates the wandering.[107] If the agent refocuses, she stops the mind-wandering. If she elaborates the mind-wandering then she transitions from this imaginative project to another, most likely directed daydreaming. Either way, these instances of mind-wandering can also be handled by the Skillful Action Account by placing their initial phases within the category of action slips.

What the account excludes are cases of mind-wandering that are (thoroughly) involuntary. Unbidden imagery can itself be a cue for some automatic mind-wandering. My stress or anxiety about writing this chapter can cause a misfire of the mechanisms underwriting my capacity for imagining, resulting in imagery of a stack of papers gathering dust in my desk. This imagery, in turn, might trigger an automatic exercise of the capacity in elaborating it into imagery of my taking the papers out of the desk or forming a thought, say, that I really better get to it. The difference between unbidden imagery and the elaboration of that imagery concerns their etiology. On the Skillful Action Account, what makes the imagery unbidden and, so, not an instance of mind-wandering is that it is not the product of an exercise of the capacity to imagine. It is merely a product of the misfiring of the mechanisms underwriting that capacity. By contrast, the latter imagery or thought counts as a constitutive means of mind-wandering because it is an exercise of the capacity. I recognize that this does not make drawing the line between unbidden imagery and automatic mind-wandering or a slip into such wandering in particular cases easy. Doing so will largely be an empirical matter. Nonetheless, the Skillful Action Account provides principled distinctions that would in principle allow such line drawing.[108]

The Skillful Action Account makes sense of instances of passive imagining by treating them as, in one way or another, noncentral. Such instances are malfunctions, non-malfunctional misfires, or action slips. The exceptions are automatic imagining and mind-wandering. The account treats the former as a key element of active imagining. Much active imagining results from the automatic triggering of the capacity to imagine in response to an appropriate cue. Indeed, the claim that imagining is paradigmatically a type of skillful action depends on the possibility of automatic imagining. Finally, the account treats mind-wandering as a subtype of

[107] This point applies equally well to at least initially undirected daydreaming. It also highlights a way in which attention is at least sometimes essential to agency: when an agent is primed to perform multiple practically conflicting actions, which action she ends up performing can at least sometimes be a matter of her attending to something. Such selection for action might consist in an agent's coming into or maintaining perceptual contact with some target or in an increase of the strength of what Springle calls a "need-input" for a particular action. I take it that such tie-breaking attention is common to humans and animals. I am not sure it's exhaustive of "attention" more generally. And I do not think that all action requires it (*cf.* Wu 2019; Bickel 2024).

[108] Thanks to an anonymous reviewer for pushing me to elaborate the Skillful Action Account's handling of mind-wandering.

(episodic) hypothetical or counterfactual thought that is necessarily triggered automatically. In that case, when it occurs, mind-wandering is either a slip or a habitual or skillful seeking-out of alternative courses of action. In all, the account's treatment of passive imagining as noncentral in relation to active imagining further bolsters and enriches its characterization of imagining.

3.6 Conclusion

According to the Skillful Action Account of Imagining, imagining is paradigmatically active and to imagine actively is to be primed to construct a discursive representation and to select its content together as a means of completing some imaginative project. Our pretender's imagining that the floor is lava just is her potentiating(/generating an intentional action acorn for) the formation of a discursive representation of the floor's being lava as a means of carrying on the pretense. And that act of imagining is in turn constituted by her practical knowledge of it as means and of her carrying on the pretense as end. The account puts forward conceptual resources for addressing the Unity-in-Multiplicity problem. By identifying imagining as paradigmatically a type of skillful action, the account can explain how imagining exhibits conflicting features and plays conflicting roles in contributing to the completion of distinct types of imaginative project. After all, it is a feature of skillful action in general that such action tolerates variation and contrary properties across its instances. What ties those instances together is the means-end order specific to the relevant type of skillful action. The account likewise puts forward the means-end order specific to imagining. The account appeals to a qualified form of intentionalism to further enrich its action-first, non-reductive characterization with an account of the intentionality of imagining. The agent imagines what she intends to imagine only when her imaginative performance matches her cogent practical cognition thereof. What the agent can successfully practically cognize, i.e, come to practically know and, thus, what she is able to imagine evolve as a function of her practicing engaging in types of imaginative project. Other things being equal, the more developed an agent's capacity to imagine, the more she is able to imagine. Finally, the Skillful Action Account treats passive instances of imagining as noncentral and holds that neither automatic imagining nor mind-wandering are genuinely passive. This treatment allows the account to maintain that active imagining is the central imaginative kind, further bolstering and enriching its characterization.

This chapter sets out a conception of imagination that is useful for philosophical theorizing about imagination. In particular, the account can serve as a basis for integrating already existing research in the philosophy of imagination. It can likewise motivate further research unifying seemingly disparate imaginative

phenomena or delineating non-imaginative phenomena from the genuine article. There are at least six avenues of for further research:

1) spelling out the guiding principles, methods, techniques, heuristics, and so on specific to each type of imaginative project;
2) accounting for types of imaginative project which crosscut those already listed, for instance, interactive fiction or thought experiment;
3) accounting for standard puzzles in the literature on imagination, for instance, how it is that our emotional responses to fictional events can be rational despite being about things we know not to be real (often called the "paradox of fiction"; Radford and Weston 1975; Walton 1978);
4) accounting for malfunctions in imagining;
5) providing a full account of the intentionality of imagining, plausibly along the lines of Berto (2017, 2018, 2023).
6) fitting the Skillful Action Account of imagining to logics of action, for instance, the see-to-it-that (stit) logic (Badura and Wansing 2021).[109]

[109] On stit logic, see Benlap, Perloff, and Xu (2001).

4.0 Remembering is an Imaginative Project

4.1 Introduction

In this chapter, I defend the claim that episodic remembering is a mental action by arguing that acts of retrieving past experiences and acts of potentiating(/generating an intentional action acorn for) the formation of discursive representations in imagination are of a kind. More specifically, episodic remembering is a type of imaginative project that involves the agent's at least potentiating the formation of some scene with the aim of (accurately) re-experiencing some specific event from her personal past. Episodic remembering, by virtue of being a type of imaginative project, is on a par with, for instance, episodic hypothetical thought, pretense, engagement with fiction, predicting or explaining others' behavior, reasoning about possibility or necessity, bare imagining for its own sake, and so on. What ties these projects together as imaginative is that their completion stands as an end to a particular means, namely, the agent's potentiated construction and selection of the appropriate representational content. Such goal-directed constructive and selective activity are the paradigmatic expression of the agent's capacity to imagine. An agent's ability to episodically remember, then, depends not just on her being able to make it the case that some memory event occurs but also on her ability to at least potentiate forming the relevant scene at will. On the view that I develop, agential scene construction as part of episodic remembering is underwritten by the agent's navigating her personal past.

My argument for the claim that episodic remembering is an imaginative project starts by drawing out two commitments of accounts of memory as a mental action (Hopkins, 2014, 2018, 2023; Arango-Muñoz and Bermúdez, 2018; Goldwasser, 2022). First, episodic remembering is relevantly like imagining in virtue of being constructive. Call this the "constructive commitment." Second, the construction involved in episodic remembering is highly constrained by how things were in the agent's past. Call this the "constraint commitment." These two commitments seem to be in tension. On the one hand, the constructive commitment seems to suggest that the agent can in principle construct a scene that does not stem from her personal past but that nonetheless counts as an instance of episodic memory (Michaelian, forthcoming). Such an implication threatens to vitiate the idea that episodic memory serves as an epistemic source by reliably connecting subjects to their past experiences. Indeed, unless there is good reason to think that paradigmatic cases of episodic memory involve retrieval of an experience from the personal past, there is no reason to think that its products can reliably justify beliefs about that past. On the other hand, the constraint commitment seems to suggest that the agent's powers for at least potentiating the construction and selection of discursive representational contents are so

limited in the context of episodic remembering as to warrant distinguishing remembering from imagining proper. Indeed, Sant'Anna (2023) argues that remembering is metaphysically distinct from imagining because the agent cannot select the content that she remembers. She cannot choose what she remembers because what she can remember is fully determined by the relevant past experience, whereas what she can imagine is in principle fully determined by her.

I suggest that the apparent tension between the constructive and constraint commitments can be resolved while holding that remembering is of a kind with imagining. I argue that the tension dissipates once we appreciate that it is the agent who at once imposes constraints on her constructive activity and satisfies those very constraints through her engagement in that activity. In fact, all rational action involves the self-imposition of some constraints on some performance and the satisfaction of those constraints through that performance. On the powers framework, the self-imposition of such constraints and their satisfaction through performance in cases of bodily[110] action amount, on the one hand, to the agent's activation of the relevant intentional action concepts and their subsequent instantiation in performance and, on the other hand, to the enormously complicated activity of her neurocognitive architecture, her musculature, things in the world under "purely physical descriptions," and the causal interactions between them.[111] What distinguishes imaginative action is that the action in question is the (potentiation of the) construction and selection of discursive representational content as a means of completing some imaginative project. And what distinguishes epistemically useful imaginative projects are the agent's self-imposition and subsequent satisfaction of specifically epistemic constraints due to her aiming to gain the relevant justified beliefs (for a similar claim made by way of a very different argument, see Proust 2013b, especially chapter 7).

For instance, in what De Brigard (2014a) calls "episodic hypothetical thought," consists in the agent's at least potentiating the construction a scene of an actual, possible, or counterfactual past, present, or future state of affairs in order to explore its causal and logical consequences (see also Kind, 2016, 2018). She will not succeed in her aim of figuring out what would (have) happen(ed) if such-and-such contingent matter of fact were the case if she does not constrain her constructive activity to stay within the confines of the relevant scenario, if she ignores or fails to hew to the relevant causal or logical relations, or if she lets her mind wander during her attempt. Given that she succeeds, what makes her imagining epistemically useful is that it is informed

[110] In keeping with chapter 2 of this dissertation, I allow that some bodily action is at once the overt performance of a mental action.

[111] Again, the "subsequent" here should be understood, in the first instance, as logical and not temporal.

by an epistemic aim, namely, to figure out what would (have) happen(ed) if such-and-such were the case. The point generalizes to other imaginative projects, including non-epistemic ones: their successful completion depends on the agentively self-constrained processes of at least potentiated construction and selection of discursive representational content. Only at the limit of bare imagining for its own sake does the agent exhibit a near complete latitude in her constructive activity with respect to the content that she represents. Nonetheless, by virtue of being an (epistemic) undertaking that takes as its means the agent's at least potentiating the construction of a discursive representation and selecting its content, episodic remembering is of a kind with other imaginative projects. In particular, episodic remembering is the agent's at least potentiating the construction of a scene of an event from the agent's actual personal past that in turn activates corresponding perceptual recognitional skills and is underwritten by her exploration of that sliver of temporal and modal space.

In what follows, I discuss the accounts of memory as a mental action and set out the constructive and constraint commitments in greater detail (Section 4.2). I then take a brief interlude into the so-called "(dis)continuism debate," (Section 4.3) from which I draw the notion of constraints as applied to imagining and remembering (Section 4.4). I argue that the notion has so far been misconstrued due to a failure to appreciate that action in general and imagining in particular are constrained and that this misconstrual is where the tension between the two commitments stems from (Section 4.5). I rectify the misconstrual and the apparent tension it creates by showing that, properly understood, episodic remembering is an epistemically useful imaginative project on a par with, e.g., episodic hypothetical thought.

4.2 Remembering as a Mental Action

Four preliminaries. First, to simplify things, going forward, I restrict focus to remembering and imagining among neurally intact adults in normal conditions. Second, following the accounts of memory as a mental action (Hopkins, 2014, 2018, 2023; Arango-Muñoz and Bermúdez, 2018; Goldwasser, 2022), I restrict focus to the retrieval of episodic memories. Instances of episodic recall are (self-)conscious attitudes or processes that contain discursive content with a sensory- or experience-like phenomenal "character" or "format," where such content represents some past event at a sufficiently far remove from the remembering subject's local temporal environment (Tulving, 1972: 385).[112] Due to its being caused in the absence of a corresponding stimulus, such

[112] As the claim that memory is a mental action focuses squarely on episodic memory (*cf*. Goldwasser, 2022), hereafter, I use "episodic memory," "episodic remembering," "episodic recall," "memory," "remembering," and "recollection" interchangeably except where clarity dictates. I also use "mnemic processing" to pick out remembering at the level of cognitive process and "mnemic projects" to pick out remembering as embodying a type of goal that mnemic processing

sensory- or experience-like content is classified as imagery (Kind 2005, 2021; Van Leeuwen, 2013; Nanay, 2010, 2016, 2021, 2023). Recalling chapters 2 and 3, on the powers framework, such content just is the activation of corresponding perceptual recognitional skills. An agent's episodically remembering her last birthday party thus consists in her being in a (self-)conscious state of entertaining imagery, say, of her being surrounded by friends, blowing out candles, eating cake, and so on, where that imagery (at least typically) matches her experience of her birthday party in the appropriate way.[113] Third, the primary imaginative kind at issue for my argument roughly corresponds to what Van Leeuwen (2013) calls "constructive imagining." Constructive imagining is, on the powers framework, the agent's at least potentiating "coming up with [discursive] representations that have [such-and-such] content" by "combin[ing] elements of ideas from memory" (221, 224). Such contents may but need not be imagistic. Importantly, recalling chapter 3, the invocation of "memory" in the definition of constructive imagining refers to the agent's conceptual repertoire as well as to her perceptual prototypes as embodied in sets of discursive solution systems for working with the same discursive contents. Such "memory" does not presuppose episodic memory. Finally, although in presenting the accounts of memory as a mental action I mark relevant differences between them, my argument does not hang on accepting any specific account. All three make the same two commitments and, in doing so, come to rely on an argument that I present at the end of this section.

Before getting to the two commitments, a brief summary of the accounts of memory as a mental action. Because all of the accounts of memory as a mental action that I will consider adopt some form of causalism, I will unpack them in their own terms. All three accounts agree that agents can episodically remember as an unmediated, non-deviant effect of intending to do so. That is, all three agree that remembering

---

aims to satisfy. Though the present essay has upshots for other forms of declarative memory, including semantic memory, I leave discussing those upshots for another occasion. Finally, "sufficiently far remove" is intentionally left vague. If one thinks one can episodically remember reading the previous sentence as one reads this one then a few seconds might be sufficiently far removed from one's local (read: present) temporal environment.

[113] Some, notably Dorothea Debus, Sarah Robins, Jordi Fernández, and Andre Sant'Anna, hold that episodic memories are (at least typically) rightly experienced by the subject as being authentic to her actual past experience (on the authenticity of memory, see Bernecker, 2010, 2015). "In the appropriate way" is meant to be neutral between causalist and non-causalist accounts of memory. On causalist accounts of memory, the match between memory and remembered experience is appropriate only if the latter non-deviantly causes the former (Martin and Deutscher, 1966). Classical causalists held that appropriate causation depended on there being stored representations of the relevant event(s), usually in the form of so-called "memory traces." Contemporary causalists about memory have developed increasingly sophisticated accounts of appropriate causation by a prior experience (for instance, Werning 2020; Langland-Hassan 2022a, 2023). On at least some non-causalist accounts like constructivism or simulationism, the match is appropriate only if it is produced by the activity of a reliable episodic construction system (De Brigard, 2014a; Michaelian 2016a,b).

is at least sometimes a so-called "basic action" (Danto, 1965; Hornsby, 1980, 2013). According to causalism about action, a basic action is one whose performance does not depend on the agent's performing some other action as a means. A classic example is raising an arm: able-bodied agents under normal conditions can intentionally raise an arm "just like that," that is, without having to do anything else to raise it. Other things being equal, on causalism, the causal efficacy of the agent's intention is sufficient. The claim, then, is that agents can intentionally remember "just like that," without having to do anything more than intend to remember.

That said, all three accounts of memory as a mental action acknowledge that there are "unbidden memories," that is, memories that come in the absence of intending to remember or against the agent's will. The claim that remembering is a mental action only requires that the paradigmatic expression of memory is intentional remembering. Unbidden memory is the exception on these accounts. Note that this is not a claim about frequency. Many or even most memories may come unbidden. Rather, the claim is that a memory's coming unbidden is its coming to consciousness as the result of something other than the proper exercise of the capacity to remember. According to the accounts of memory as a mental action, the proper exercise of the capacity is the exercise of mnemic agency in acts of remembering. The idea is that intentional remembering takes metaphysical priority over unbidden memory.

The accounts of memory as a mental action differ on how much influence the agent has over the process of retrieval. I go in order from least to most influence. Starting with Hopkins (2014), he holds that the agency involved in remembering is exhausted by our trying to remember and either failing or succeeding (314-315, 323-325). For Hopkins, remembering is necessarily basic and what is remembered is completely "causally controlled by how [the past event] was" (324). What it means for the past to causally control what is remembered is unclear. However, it is clear that, according to Hopkins, at least the *occurrence* of a memory event is within the agent's control. Arango-Muñoz and Bermúdez (2018) hold that our mnemic agency is exhibited in the metacognitive feelings of, for instance, fluency, recognition, or familiarity that are directed at and attenuate the construction process involved in episodic remembering (84ff.). Exercise of that agency in the reconstruction of scenes from the personal past is always indirectly mediated by these feelings and the agent's practical responses to them. According to Arango-Muñoz and Bermúdez, then, the occurrence of a memory event is up to the agent and what she remembers is under her control indirectly. The content of memory is determined by the interplay of the first-order activity of construction and the second-order metacognitive feelings by which the agent then indirectly guides that first-order activity. Finally, I hold that the type of agency involved in remembering is skillful (2022: 6ff.) Because successful mnemic activity is often habitual, direct intervention by

the agent is usually unnecessary. However, at the time at least, I held that agents can directly intervene in the process of construction when mnemic activity threatens to go awry. That is, both the occurrence of a memory event and what the agent remembers are in the first instance direct products of the exercise of her mnemic agency. Now I would say that agents can directly intervene on potentiations of acts of discursive representing constitutive of covert episodic remembering and on fully performed acts of discursive representing constitutive of overt episodic remembering.

Now, onto the two commitments. First, accounts of memory as a mental action are committed to claim that remembering is like imagining in a way relevant for agency by virtue of being constructive. Call this "the constructive commitment." Second, such accounts are also committed to the claim that what an agent remembers is constrained by her past experiences. Call this "the constraint commitment." Starting with the constraint commitment, despite variation in the amount of influence the agent has over her remembering, all three accounts hold that remembering is successful (typically) only when the content constructed (sufficiently) matches the corresponding past experience in the appropriate way. How each account cashes out this requirement differs. Hopkins (2014) cashes out the constraint commitment in terms of causal control by the past event. Arango-Muñoz and Bermúdez (2018) do so in terms of the successful deployment of accuracy-oriented metacognitive feelings to correct or discard the products of the first-order activity of the relevant constructive mechanisms. And I (2022) cashed out the commitment in terms of the skillful construction of scenes appropriately matching those of past experiences, which accuracy-oriented constructive activity agents are trained to do and practice. On all three accounts, failing to represent an experience from the agent's personal past (typically) disqualifies the instance of construction in question from counting as an instance of remembering. In cases of failure, the agent misremembers, confabulates, or makes some other memory error. The claim that an agent remembers (typically) only if the content of what she constructs (sufficiently) matches some past experience in the appropriate way implies that all three accounts commit to memory's being constrained by the personal past.

Moving on to the constructive commitment, note that this commitment is not just about imagery. It is uncontentious that episodic remembering and some forms of imagining both trade in imagery (Langland-Hassan, 2015; Liao and Gendler, 2019; Kind, 2001, 2021; Nanay, 2010, 2016, 2021, 2023). What is contentious is whether the constructive processes involved in, say, forming experiential imagery of merely counterfactual or hypothetical scenes are the same in kind as those involved in recalling events from the personal past. In sharing the constructive commitment, accounts of memory as a mental action share a commitment to an identity at the

level of process between remembering and at least some forms of imagining. In particular, all three accounts seem to agree that at least intentional episodic remembering and intentional episodic hypothetical thought are both epistemic kinds of agential scene construction.

The constructive commitment, then, consists in the claim that episodic remembering and (at least some forms of) imagining are at bottom processes of agential scene construction. For evidence of the constructive commitment, one need look no further than the title of Hopkins's (2018): "Imagining the Past." Arango-Muñoz and Bermúdez (2018, 78 *passim*), like Hopkins (2014, 318-319), appeal to the control we enjoy (or lack) over conjuring mental imagery of such things as pink elephants as reason to think that we enjoy a similar degree of control over remembering. They say outright, "we assume the imaginative reconstructive conception [of episodic remembering]" (2018: 75). Likewise, in my (2022), I assumed that the imaginative reconstruction of scenes involved in remembering is often enough itself the skillful work of the remembering agent. Part of my argument here is to substantiate that assumption with the account of imagining as a skillful type of mental action that I presented in chapter 3.

The constraint commitment is adapted from philosophical and psychological theorizing about memory more broadly. According to such theorizing, for memory to serve as an epistemic source concerning past events and for it to serve its various psychological functions, its exercise must at least typically result in representations appropriately constrained by the remembering agent's past experience. By adopting a broadly naturalistic approach to memory, then, the accounts of memory as a mental action can accept some form of the constraint commitment for free. If nothing else, the accounts should accept the constraint commitment so as to avoid speaking past other philosophers of memory or memory scientists. By contrast, the constructive commitment is not so easy to justify. The accounts of memory as a mental action seem to have something like the following argument in mind in holding it:

The Argument from the Constructive Nature of Imagination

P1) Imagining is active in that, other things being equal, the agent imagines when she intends to imagine and imagines what she intends to imagine.

P2) If remembering is the same kind of agential activity as imagining then, other things being equal, the agent remembers when she intends to remember and remembers what she intends to remember.

P3) Remembering is the same kind of agential activity as imagining.

C:) Other things being equal, the agent remembers when she intends to remember and remembers what she intends to remember.

Recall from the previous chapter that (P1) is a widely accepted claim about imagination that goes at least as far back as Hume (1777/2000, 1739/2007) and has been aptly labeled the "the freedom of imagination" (for an articulation of this claim that distinguishes occurrence and content, see Dorsch, 2012, especially chapters 13 and 14; *cf.* Strawson 2003). What is important about this premise in the context of my argument here is that the agent's freedom to select what she imagines is taken to distinguish imagination as a faculty from perception, belief, and so on. (P2) is, arguably, a presupposition of both parties to the debate on whether remembering is in fact a mental action. In adopting the constructive commitment, proponents commit to (P3) and this only makes sense if they think that remembering's being like imagining in the relevant respect allows them to draw the conclusion that remembering is a mental action. With the exception of Strawson (2003), opponents will allow that imagining is paradigmatically active. In which case, they ought to accept (P2) as vacuously true. (P3) is supposed to be substantiated by adopting the constructive commitment.

Note that, as formulated, the Argument from the Constructive Nature of Imagination risks creating a tension between the constructive commitment and the constraint commitment. The freedom of imagination consists not just in the agent's ability to bring about occurrences of imagining by intending to imagine. Otherwise, imagination would arguably be no different from judgment or paying attention. Other things being equal, to bring about the occurrence of a judgment, one need only rehearse a judgment that one is already committed to or make a novel trivial judgment, say, that human toenails do not sprout wings. Similarly, other things being equal, to pay attention, one need only intentionally perceptually fix on something arbitrarily. By contrast, the agent's ability to select the content of a (potentiated) discursive representation is at least partly definitional of imagining (Dorsch 2012). But if what makes remembering a mental action is its being like imagining in those respects that distinguish the latter from other mental faculties then it seems that remembering should also admit of the agent's selection of content. In that case, it seems that remembering should turn out unconstrained by the past. Yet, the accounts of memory as a mental action share the constraint commitment. And in that case, the content of memory seems not to be something that the agent can select. It then becomes difficult to see how remembering could consist in the same agential activity as imagining. And this brings us full circle. It is the primary aim of this chapter to resolve the apparent tension between the constructive and constraint commitments. I do so by substantiating (P3) in a way that shows that all imagining is subject to the agent's selection of content while nevertheless being similarly constrained. In particular, I argue that all rational action is self-constrained, that some constraints on imagining are specifically epistemic, and that, of these constraints, those placed on (potentiating) the construction of scenes that appropriately match the

actual or nearby possible past, present, or future are of a kind. However, before getting there, a brief interlude on the so-called "(dis)continuism debate."

4.3 A Brief Interlude Into the (Dis)continuism Debate

Endorsing the constructive commitment puts the accounts of memory as a mental action on one side of an ongoing debate in contemporary philosophy of memory, namely, the (dis)continuism debate (for thorough overviews, see Perrin, 2016; Michaelian, 2016a; Perrin and Michaelian, 2017; and Michaelian, Perrin, and Sant'Anna, 2020). At issue is whether episodic remembering is of a kind epistemically or metaphysically with imagining, specifically episodic hypothetical thought (De Brigard 2014a). According to the literature in the (dis)continuism debate, episodic hypothetical thought consists in constructing an episodic representation of a possible or counterfactual past, present, or future state of affairs in order to explore its causal and logical consequences. The (dis)continuism debate thus concerns whether remembering has the same epistemic upshots as or even constitutes a dimension of such thought.

My focus will be on the metaphysical (dis)continuism debate. In this part of the debate, being of the same kind means being determinants of the same determinable process, attitude, or both. Continuists claim that remembering and other forms of episodic thought are metaphysically of a kind as processes, attitudes, or both (see, for instance, De Brigard, 2014a; Michaelian, 2016a,b, 2021, 2022; Langland-Hassan, 2021, 2022a,b,c, 2023). By contrast, discontinuists claim that remembering is metaphysically distinct from other forms of episodic thought (see, for instance, Perrin, 2016; Robins, 2020; 2022; Sant'Anna, 2023). Accordingly, the process of remembering is not of a kind with that of episodic hypothetical thought or is a distinct kind of attitude from the kind produced by such thought. The metaphysical (dis)continuism debate admits of extreme and moderate positions and allows for hybrid views. For instance, one can be a process continuist and an attitude discontinuist or a process discontinuist and an attitude continuist.[114] By endorsing the constructive commitment with a view to substantiating (P3) of the Argument from the Constructive Nature of Imagination, accounts of memory as a mental action at minimum commit to a form of metaphysical process continuism.

However, the route by which the accounts of memory as a mental action arrive at process continuism differs significantly from how metaphysical continuists have tended to defend the position. Indeed, until recently, the metaphysical (dis)continuism debate has focused primarily on appealing to empirical

[114] To my knowledge, no one has yet argued for the latter kind of hybrid view.

considerations. For instance, continuists appeal to substantial overlap in the neural correlates of construction and those of remembering (Klein et al., 2002; Rosenbaum et al., 2005; Addis et al., 2007; Schacter and Addis, 2007a,b; Szpunar et al., 2007; De Brigard et al., 2013; Schacter et al., 2015). Discontinuists argue that such overlap underdetermines the metaphysical nature of these processes on the grounds that the relevant neural correlates appear to be at work in much more than scene construction and remembering (D'Argembeau and Van der Linden, 2004; Schacter and Addis, 2007a; Addis et al., 2007; Szpunar et al., 2007; McDonough and Gallo, 2010; Martin et al., 2011). However, it is open to the continuist to argue that the overlap is plausibly more local than it might have at first seemed or involves patterns of activation specific to construction in episodic thought in general.

Similarly, continuists claim that there are striking similarities between the cognitive and biological functions of remembering and episodic hypothetical thought. In particular, remembering and other forms of episodic thought appear to comprise a single episodic construction system. This system has the function to simulate events along a number of temporal and modal dimensions: past, present, or future on the one hand and possible, actual, counterfactual, or necessary on the other (Bartlett, 1932; Hassabis and Maguire, 2009; De Brigard, 2014a; Michaelian 2016a,b; Wagoner, 2017). In other words, according to continuists, remembering constitutes one highly constrained dimension of a more general faculty for "mental time travel." Discontinuists reply that there are significant functional differences between remembering and episodic hypothetical thought. For instance, remembering appears to have a distinct functional role by virtue of requiring the activation of so-called "memory traces." Memory traces are a theoretical posit that is supposed to explain how experiences of events are encoded such that they can later be retrieved (Martin and Deutscher, 1966; Debus, 2008, 2014, 2016; De Brigard, 2014a; Klein, 2016; Perrin, 2016; Fernández, 2017, 2019; *cf.* De Brigard 2014b). There is some empirical evidence that such traces exist in the neurocognitive architecture. And activation of traces for representing particular events *appears* to be specific to the well-functioning of memory. However, it is open to the continuist to point out both that traces are likely distributed and that, *in fact*, they implicated in the well-functioning of episodic hypothetical thought as well (Hoffman and McNaughton, 2002; Irish and Piguet, 2013; Robins, 2016; Wagner, 2016; Aronowitz, forthcoming).

Finally, continuists argue that phenomenological differences between remembering and imagining, like the former's exhibiting a "feeling of pastness" (Tulving, 1983, 1985; Robins, 2020), can be attributed to metacognitive mechanisms discussed in the empirical literature under the heading of "metamemory" (Dunlosky and Bjork, 2008; Dunlosky and Tauber, 2016). In that case, first-order mnemic processing might well be the

same as that involved in episodic hypothetical thought. Discontinuists argue that such phenomenology is at least partly definitional of remembering as a personal-level attitude (D'Argembeau and Van der Linden, 2004, 2006; Addis et al., 2010; Berntsen and Bohn, 2010; De Brigard and Giovanello, 2012; Rasmussen and Berntsen, 2013; Robins, 2020).[115] And discontinuists might well argue that such phenomenology is not due to metacognition but to the first-order processes involved specifically in remembering. However, it is open to the continuist to argue that the jury is still out on whether the cause of memory-specific phenomenology is due to first-order processes rather than metacognition.

Empirical considerations alone have only moved the metaphysical (dis)continuism debate so far. Indeed, the debate appears to be in gridlock. Moreover, the Argument from the Constructive Nature of Imagination that is endorsed by accounts of memory as a mental action does not make reference to any such empirical considerations. Granted, one could attempt to substantiate (P3) of the argument by appeal to those considerations, as Arango-Muñoz and Bermúdez appear to do implicitly in endorsing what they call "the imaginative reconstructive conception" of episodic remembering. However, I follow Langland-Hassan (2022a, 2023) and Sant'Anna (2023) in thinking that we can make progress in the metaphysical (dis)continuism debate by shifting focus away from empirical considerations and towards metaphysical ones. In particular, I suggest taking a closer look at the operation of memory and imagination within the context of the structure of mental agency. I believe that progress can be made by considering the role that constraints play in the processes involved in remembering and imagining. It is to this that I now turn.

4.4 The Role of Constraints in Memory and Imagination, Part I

4.4.1 Constraints

So far, my discussion of constraints has been limited to the constraint commitment, the claim that what an agent remembers is constrained by her past experience. In considering Langland-Hassan's (2022a, 2023) and Sant'Anna's (2023) recent work on the metaphysical (dis)continuism debate in this section, I broaden the discussion to include constraints supposedly imposed on other forms of imagining. A burgeoning literature on the possible epistemic uses of imagination takes constraints as its focus. Philosophers working in this literature consider whether imagination is constrained and, if so, whether it is constrained in ways that allow agents to form justified beliefs on the basis of their imagining (see, for instance, Dorsch, 2012; Kind and Kung (eds),

[115] Again, see Drayson (2014) on the personal/sub-personal distinction.

2016; Kind and Badura (eds), 2021; Myers 2021a,b, 2024, forthcoming a,b). For instance, in attempting to figure out whether her sofa will fit through a particular door frame, an agent might visualize the couch going through it. [116] For her visualizing to justify her belief that the sofa will (not) fit, the imagery that she forms must be constrained, on the one hand, by how the world actually is and, on the other hand, by what would causally or logically follow from the world being as she imagines it. Specifically, the imagined spatial dimensions of the sofa and of the door frame must match their actual spatial dimensions such that the agent's imagining their hypothetical interaction tracks to a sufficient degree of accuracy what would result from their actual interaction.

As I understand them, constraints are rules to which an activity adheres or limitations applied to a process. As the sofa visualization case suggests, the imposition of constraints is partly definitional of the activities they are imposed on by giving those activities a determinate shape. The agent visualizing the sofa is not engaged in an act of figuring something out absent their imposition. Yet, constraints are not fully definitional of the corresponding activities, since those activities themselves are what satisfy the relevant constraints by taking the shape set out by the imposition of those constraints. In the sofa visualization case, the imposition of the reality-oriented constraint on the agent's imagining sets out the determinate shape that her imagining is to take in light of her project, namely, to figure out whether her sofa will fit through the door frame. But it is her successfully visualizing the sofa and door frame interacting with their actual spatial dimensions that realizes her aim and, in so doing, allows her to complete her project. Speaking somewhat metaphorically, constraints form a mold for the activities that they are imposed on and those activities are cast into that mold.

Constraints can be lateral or top-down (Langland-Hassan 2016). Lateral constraints are rules or limits imposed in the first instance by things other than the agent. That is, such constraints are imposed by the agent's neurocognitive architecture or other aspects of her biology or by aspects of the world outside of her control. For instance, a lateral constraint on experiencing the sensation of spiciness is activation of TRPV1 receptors in taste receptor cells. Experiencing the sensation of spiciness is not something an agent can simply decide or choose. Rather, her having such an experience depends on the constitution of her taste receptor cells. By contrast, top-down constraints are imposed by the agent. A top-down constraint on experiencing the sensation of spiciness on particular occasions is ingesting foods containing capsaicin. Under normal conditions, an agent imposes this

[116] Though my examples of active imagistic imagining will primarily evoke visual imagery, my argument does not depend on imagery being of or being like that of any particular sense modality. On sensory and experiential imagination, see Kind (2001, 2021), Dorsch (2012, 37-40, 44-45), Hopkins (2024).

constraint in intending to eat something spicy and her intentionally eating some spicy thing satisfies that constraint.

The satisfaction or imposition of at least some lateral constraints is affected by the imposition of some top-down constraints both on specific occasions and over time. Satisfaction of the lateral constraint of TRPV1 receptor activation on experiencing the sensation of spiciness can be modulated on a specific occasion, say, by the choice to forgo eating something spicy or by ingesting a spoonful of yogurt before ingesting something containing capsaicin. In either case, a lack of TRPV1 activation or its being masked means that the lateral constraint on experiencing the sensation of spiciness goes unsatisfied. Similarly, imposition of the TRPV1 activation constraint can be modulated over time by repeatedly ingesting capsaicin. One builds a tolerance to spicy food by repeatedly eating it. If an agent is building such a tolerance then she will need to ingest more capsaicin to activate her TRPV1 receptors than an agent who altogether avoids spicy foods. The former's eating habits results in her experience of the sensation of spiciness being more highly laterally constrained, requiring her to ingest increasing amounts of capsaicin.

Returning to memory and imagination, assuming that imagination is appropriately constrained at least in its epistemic uses, the question arises whether remembering is a kind of epistemically useful imagining. Both Langland-Hassan (2022a, 2023) and Sant'Anna (2023) take up this question directly and arrive at opposite conclusions. Langland-Hassan answers in the positive: remembering is the agent's constructively (and accurately) imagining scenes from specific events in her personal past using "props" derived from multiple distributed memory traces. Sant'Anna answers in the negative: the lateral constraint on remembering that its content appropriately match the corresponding specific past experience cannot be satisfied by a process of construction where the agent can select the content she constructs. Call the lateral constraint on remembering that its content appropriately match the corresponding specific past experience "the specific-past-experience constraint." In the rest of this section, I draw out each philosopher's conception of the specific-past-experience constraint as well as their conception of how it is imposed and satisfied. Doing this is instructive for two reasons. First, the way that Langland-Hassan and Sant'Anna each conceive of the constraint is widespread among both philosophers of imagination and philosophers of memory. So, understanding their explications of it is helpful for understanding the standard conception. Second, it will turn out that the apparent tension between the constructive and constraint commitments stems from this conception of the constraint. So, to have any hope of resolving that apparent tension and securing the conclusion of the Argument from the

Constructive Nature of Imagination, we must adopt a conception of the specific-past-experience constraint that differs in the appropriate ways from the one adopted by Langland-Hassan and Sant'Anna.

4.4.2 Langland-Hassan's Disjunctive, Prop-Based, Causal Constructive Theory of Memory

Langland-Hassan (2022a) argues in favor of metaphysical process continuism on the grounds that the constraints imposed by memory traces and satisfied by their activation very likely crosscut scene construction in remembering and episodic hypothetical thought.[117] According to Langland-Hassan, our best empirical evidence, while not conclusive, suggests that memory traces function as "props" rather than as full-blown representations of specific past events (25-30, fn.6; see also Robins, 2016; Aronowitz, forthcoming). According to what he calls the "prop theory of traces," rather than encoding past experiences as isolable wholes, memory traces encode more generic scene elements that can be activated in representing similar past experiences as well as counterfactual or possible episodes. For instance, in remembering eating salmon for dinner last night, imagining having had halibut instead, imagining eating mahi-mahi right now, and imagining having swordfish for dinner tonight, the agent may well activate the same memory trace(s) for their dining room. The lateral constraint that the appropriate memory traces be activated, then, is not specific to memory. Appropriate trace activation is required for successful episodic thought in general. Continuity of process comes with the claim that trace activation is part of scene construction in all three temporal dimensions and all modal dimensions of mental time travel. Such activation does not distinguish the process of remembering from episodic hypothetical thinking.

This leaves the specific-past-experience constraint so far untouched. Langland-Hassan seems to be sensitive to this. Indeed, he points out that the prop theory of traces and reconstructive accounts of episodic remembering both fail to "offer an obvious means for distinguishing successful remembering from (merely) accurately representing one's personal past" and fail to "distinguish remembering from accurate representation via *deviant* causal chains such as [...] cases of relearning" (2022a: 30; original emphasis). To account for the imposition and satisfaction of the specific-past-experience constraint, Langland-Hassan appends reconstructive accounts of episodic remembering with a disjunctive causal condition. He holds that imposition of the specific-past-experience constraint ultimately stems from the fact that memories have to be appropriately caused. However, unlike classical causalists about memory (see footnote 112), he denies that there is a single causal route to satisfying the constraint.

[117] I do not focus on Langland-Hassan's (2023) because it makes roughly the same argument.

Instead, the idea is that acts of remembering are complex affairs that involve more than just trace activation. Successful remembering often also involves the application of background beliefs as well as metacognitive source-monitoring and feelings of fluency (31). These causes are neither individually sufficient nor in every case required to satisfy the specific-past-experience constraint. Sometimes the right background belief carries the day, while other times high fluency or accurate source-monitoring does, while still other times only the right combination of causes results in the agent's remembering. The thought, then, is that an agent remembers just in case she constructs a scene using trace-based props that accurately represents a past experience of hers, where this construction is partly an effect of some disjunctive set of appropriate causes. Langland-Hassan holds that identifying the members of the disjunctive set is "an empirical project" and assures us that "there will indeed be some limited number of [such causes]" (32).

It is the disjunctive causal condition of Langland-Hassan's account that is supposed to explain the satisfaction of the specific-past-experience constraint. On this way of conceiving of it, the constraint is satisfied by a motley of different activities within the neurocognitive architecture. The imposition of the constraint is due to the epistemic and psychological roles that, at least according to philosophical and empirical theorizing, memory is supposed to play. Langland-Hassan's conception of the specific-past-experience constraint thus locates both its imposition and satisfaction in aspects of the remembering agent that are outside of her control. The epistemic and psychological functions of memory are cognitive and biological matters of fact. Moreover, the activation of traces in scene construction and the operation of metacognitive mechanisms like fluency processing occur entirely outside of conscious awareness. And while the contribution of background beliefs can be conscious, say, in cases of (self-)conscious inference, their contribution is due to relations of (inferential) justification that hold independently of the agent's intentions or desires.

Summing up, Langland-Hassan appears to endorse the constructive commitment. Moreover, in line with the Argument from the Constructive Nature of Imagination, he argues for a form of metaphysical process continuism in light of this commitment. Finally, in motivating the disjunctive causal condition, he endorses the constraint commitment as well. Taken together, then, Langland-Hassan holds that remembering is an epistemically useful type of constructive imagining under constraints. However, the constraints specific to remembering—in particular, the specific-past-experience constraint—are supposed to be imposed and satisfied in light of the activity of various parts of the agent's neurocognitive architecture that lie outside of her control. Moreover, such constraints are only satisfied on a given occasion by the functioning of *some* of the relevant bits of that architecture.

4.4.3 Sant'Anna's Argument from the (In)ability to Select Content

Sant'Anna (2023), in response to Langland-Hassan (2021), argues in favor of metaphysical process discontinuism on the grounds that the constraints imposed on remembering differ from those imposed on imagining. Specifically, only in imagining can the agent deliberately choose the content that she represents. Sant'Anna qualifies his position by allowing that agents can exhibit a considerable amount of control over remembering. Agents can intentionally initiate or terminate retrieval, fix the general subject matter of their remembering, leave out some details of the remembered experience in favor of focusing on others, and even play with the temporal order of the events remembered. With respect to initiation, termination, subject matter, inclusion of details, and temporal order, remembering is relevantly similar to imagining in that the agent enjoys control over these aspects of her acts of remembering.

Sant'Anna claims that the discontinuity between remembering and imagining shows up when we consider an agent's (in)ability to intervene deliberately on scene construction. As he puts it:

> [I]t is not the case that we have control over the contents of remembering. Rather, the information that figures in remembering is *given* to us at the time of retrieval, and, more importantly, it cannot be altered by us after it has been retrieved (2023: 11; original emphasis).
>
> The suggestion is that we cannot alter the content of remembering by consciously including information that is not available for retrieval in representing a subject matter (2023: 9).
>
> [R]emembering is constrained by a source evaluation that the information entertained in mind originates in the past [whereas imagining is not]. And since it is not under our control that remembering involves such a source evaluation, we experience remembering as being necessarily constrained in this way (2023: 17).

Paraphrasing, the idea is that imposition of the specific-past-experience constraint for some reason precludes the agent from knowingly injecting novel content into an act of remembering. By contrast, the agent can inject whatever content she likes into an act of imagining without thereby failing to imagine. This asymmetry between remembering and imagining with respect to the agent's ability to select the content she represents suggests that there is some deep dissimilarity between them as processes.

To locate the root of the asymmetry, Sant'Anna draws attention to the role of source-monitoring in remembering. Source-monitoring is a function of metacognition. As was noted already, metacognition is among the features of the neurocognitive architecture that contributes to satisfying lateral constraints like the specific-past-experience constraint. Source-monitoring plays a role in both remembering and imagining. Well-functioning source-monitoring in imagining delivers evaluations that what is imagined originates in the imaginer's mind rather than in the world. By contrast, well-functioning source-monitoring in remembering

delivers evaluations that what is remembered originates in the world rather than in the rememberer's mind. It is source-monitoring of a constructed scene for its origination in a specific past experience that precludes the agent from injecting novel content into remembering. Assuming her metacognitive machinery is intact, she cannot freely inject whatever novel content she likes and, yet, continue to treat the resulting scene as originating in the world on pain of lapsing into irrationality. By contrast, so long as some of the content of an agent's imagining is rightly evaluated as originating in her mind, she can freely inject stored content from memory into her imagining without cost. Having located the root of the asymmetry between the agent's ability to select the content of her imagining and her inability to select the content of her remembering, Sant'Anna draws the conclusion that remembering and imagining are metaphysically distinct kinds of processes.

Note that Sant'Anna's argument depends on the same conception of how the specific-past-experience constraint is imposed and satisfied as Langland-Hassan's. Both hold that the constraint is imposed as a function of cognitive and biological matters of fact. And both hold that the constraint is satisfied by the well-functioning of the relevant parts of the agent's neurocognitive architecture. Langland-Hassan thinks that holding both claims is consistent with holding in addition the claim that remembering and constructive imagining are at bottom the same kind of process. For his disjunctive causal condition denies that activity in each of the relevant bits of the agent's cognitive architecture is required for successful remembering. It is not a requirement on this condition that every possibly appropriate sufficient cause of remembering function as a cause in each instance of remembering. Nonetheless, presumably, with Sant'Anna, he would deny that the agent can inject whatever novel content she wants into scene construction and still count as remembering. For Langland-Hassan, there might be cases where the agent can get by without source-monitoring. But unless those cases somehow turn out to constitute the bulk of remembering, the continuist part of his causal constructive account appears to rest on extremely shaky ground.

In effect, what Sant'Anna points out is that conceiving of the imposition and satisfaction of the specific-past-experience constraint as both he and Langland-Hassan do requires holding that the agent (at least typically) lacks a kind of freedom with respect to remembering that she seems to (always) enjoy with respect to imagining. Namely, the agent is not truly free to remember *what* she intends to remember because the well-functioning of a bit of her neurocognitive architecture fixes what she remembers on any occasion of successful active remembering. In that case, appeals to the shared constructive nature of imagination and memory by accounts of memory as a mental action seem unwarranted. Yet, both Langland-Hassan and Sant'Anna seem to want to uphold the claim that remembering is a mental action in part because of the role that the agent has in

scene construction. In that case, the role of source-monitoring as what underlies the need to distinguish remembering and imagining as processes seems overblown. The apparent tension between the constructive and constraint commitments reemerges here.

4.5 The Role of Constraints in Memory and Imagination, Part II

We have seen that conceiving of the specific-past-experience constraint as both imposed and satisfied ultimately by things outside of the agent's control leads to the apparent tension between the constructive and constraint commitments. At this point, one might be tempted to give up on the constructive commitment and accept Sant'Anna's negative argument. Due to the well-functioning of source-monitoring, the agent cannot get creative with the content of her memory in the way that she can with the content of her imagining. It is this difference that supposedly makes for a metaphysical difference between the process of remembering and that of imagining. One might think, then, that we ought to deny (P3) of the Argument from the Constructive Nature of Imagination presented in Section 4.2 and draw the negation of the antecedent of the conditional forming (P2), namely, that it is not the case that remembering is the same kind of agential activity as imagining. And, in that case, one might well wonder why the accounts of memory as a mental action seem to endorse the argument. Maybe remembering can be a mental action without being relevantly like imagining.

In response, I want to suggest that this line of thought is mistaken. In particular, I want to suggest that proponents of the claim that remembering is a mental action should hold that, properly understood, remembering is a type of epistemically useful imaginative project. Imagining is active at least in part due to its involving the agent's self-imposing and subsequently satisfying constraints. Such activity is imaginative in that it consists in the agent's at least potentiating(/generating an intentional action acorn for) the construction of some discursive representation and selecting that representation's content as a means to completing some imaginative project. The occurrence and content of any imagining are therefore constrained to some degree. With respect to epistemically useful imaginative projects, the relevant constraints are imposed with a view to forming veridical imagery or justified imaginings and are satisfied by the agent's forming such imagery or imaginings. And with respect to episodic remembering and episodic hypothetical thought, the relevant constraints are imposed by a limit on navigating modal and temporal space to some circumscribed region thereof. These constraints are satisfied by the agent's at least potentiating the construction of scenes underwritten by the navigation of the appropriate part of modal and temporal space. Recalling the previous chapter, the aim specific to episodic remembering is re-experiencing the event and is underwritten by navigation to the relevant sliver of modal and

temporal space as what McDowell (2022) calls a "suitable happening" (220-221). Thus, my suggestion is that proponents of the claim that remembering is a mental action substantiate (P3) of the Argument from the Constructive Nature of Imagination by appealing to a conception of the specific-past-experience constraint that differs significantly from the one endorsed by Langland-Hassan (2022a, 2023) and Sant'Anna (2023). Specifically, proponents should appeal to a conception of the specific-past-experience constraint that treats it as being ultimately imposed by the agent's aim in engaging in acts of episodically remembering and as being satisfied by her successful (potentiation of the) construction of the relevant scene sufficient to activate corresponding perceptual recognitional skills, underwritten by navigation of her personal past.

4.5.1 Self-Determination as the Self-Imposition and Satisfaction of Constraints

I begin by recalling from chapter 1 what I take to be an essential feature of intentional agency and what I there described as the agent's structuring her own agency. Here, I call this feature "self-determination" and elaborate it in part by describing it in terms of constraints. The idea is this: when performing an intentional action, the agent does two things. First, in intending to do *A*, the agent self-imposes constraints on her token doing *A*. For instance, in intending to raise her arm, the agent determines that her arm rises rather than lowers and, thus, that her subsequent movement has properties that characterize it as an arm raising. But this is not all that she does. The agent also fixes those properties that in the context of action satisfy those constraints and thereby come to constitute her performance. In intentionally raising her arm, the agent not only determines that her arm rises but properties of its trajectory in accordance with the constraints that she has imposed in intending to raise her arm. Taken together, the agent's determining her action is her imposing constraints on her non-mental or mental movement and subsequently satisfying those very constraints through making the relevant movements.

Importantly, the distinction between the imposition of the relevant constraints and their satisfaction as distinct parts of the agent's self-determination is conceptual rather than metaphysical or temporal. The "subsequently" just used is first and foremost logical even when the agent's movements occur later in time from when she begins imposing constraints on those movements. Moreover, imposition and satisfaction as part of the agent's determining her action are often metaphysically and temporally coincident, specifically when she acts spontaneously. In cases where the agent intentionally raises her arm without previously deliberating, her intentional action at once imposes and satisfies the relevant constraints on her movement. Self-determination is an essential feature of intentional agency in that such agency is of its essence self-guided movement (Frankfurt, 1978; for similar claims applied to mental action on very different grounds, see Wu 2011a,b, 2013a,b, 2014,

2016, 2019, 2022, 2023b).[118] With respect to action at least, I want to replace the previous metaphor of mold casting and filling with this notion of self-determination.

Most importantly, the constraints imposed by the agent in determining her action include lateral constraints on her performance. Recall that, by definition, such constraints are not imposed in the first instance by the agent. Yet, recall as well that the agent can (at least) modulate, mask, and indirectly impose such constraints. That her arm movement exhibits limb dynamics appropriate for its rising is a lateral constraint on that movement that limits the arm's available limb dynamics. It is in the first instance the agent's motor system that fixes the parameters of the limb dynamics of her arm raising, not her. She cannot directly fix these parameters herself (Butterfill and Sinigaglia 2014; Mylopoulos and Pacherie 2016; Burnston 2017; Fridland 2021). Nevertheless, it is the agent that ultimately imposes this constraint, albeit indirectly through imposing the top-down constraint that she raise her arm.[119] She could free herself from the imposition of the lateral constraint by abandoning her intention to raise her arm. And she can modulate its imposition, say, by respecifying her intention to raise her arm in such-and-such a way. She can even decide mid-movement to adjust her arm raising and this will modulate the imposition of the constraint that her arm exhibit some appropriate set of limb dynamics. In either case, the result is a decrease in the members of the set of appropriate limb dynamics and, conversely, an increase in the stringency with which the lateral constraint is imposed.

Another way to put the point is that self-determination concerns the way(s) in which an action is performed. The imposition of both top-down and lateral constraints as well as their subsequent satisfaction in an instance of doing *A* are the agent's applying her knowledge-how of some way of doing *A* through self-determination and in light of her intention (on knowledge-how, see Fantl, 2008; Bengson and Moffett, 2011; Carter and Poston, 2018; and Pavese, 2023). This way of thinking about self-determination applies to basic action. Even on the powers framework, raising an arm (as part of a constitutive unity, say, of raising an arm to cast a vote) can be basic and, yet, admits of being done in multiple ways.[120] Any instance of intentional arm

[118] Note: I am not here giving a full account of intentional agency. Rather, I am pointing to one structural feature of it that I take to be exhibited in instances of intentional episodic remembering and intentional episodic hypothetical thinking. See O'Brien (2017) for some further explication of this feature.

[119] The way that I am using, on the one hand, "in the first instance" and, on the other, "ultimately" or "at bottom" is reminiscent of the distinction between triggering and structuring causes (Dretske, 1988) as well as the distinction between proximal and ultimate causes (Mayr, 1961; Aaby and Ramsey, 2022). Nonetheless, because I rejected causalism in chapters 1 and 2, I want to resist wholesale adoption of either distinction here.

[120] A limiting case is one in which there is only one way for the agent to do *A*. In such a case, her determining what her doing *A* consists in collapses into her determining that she does *A* at all. The range of tracks for the manifestation of her

raising involves the agent's determining some such way, simultaneously imposing and satisfying the relevant top-down and lateral constraints on raising her arm in that way. And because the performance of non-basic actions ultimately bottoms out somewhere in the performance of some basic action, self-determination extends from the latter type of action to the former. Assuming this is right, self-determination turns out to be an essential feature of intentional agency. Action as self-guided movement generally involves the agent's self-imposing and satisfying both top-down and lateral constraints on her performance.

4.5.2 Self-Determination in Memory and Imagination

Returning to the freedom of imagination, given that self-determination is an essential feature of intentional agency generally, we should expect an agent's acts of imagining to exhibit her self-imposing and satisfying the appropriate top-down and lateral constraints. In the previous chapter, I argued that active imagining is the paradigmatic imaginative kind and proposed a view on which active imagining consists in the agent's at least potentiating the construction of some discursive representation and selecting its content as a means to performing some further learned behavior like pretending, engaging in fiction, predicting or explaining others' behavior, reasoning about possibility and necessity, episodic hypothetical thought, and, at the limit, imagining for its own sake. That active imagining is partly defined by its role *as a means* implies that any instance of active imagining is constrained by the end for which that instance serves as a means. That active imagining is partly defined by this role likewise implies that each instance that succeeds thereby satisfies the relevant constraint(s). Agents therefore determine their active imagining on this view. I have not discussed constraints anywhere else in this dissertation. So, I summarize the account of imagining I provided in the previous chapter and discuss the upshots for the notion of constraints as well as for thinking of episodic remembering as a type of imaginative project.

On the account that I propose, imagination as a faculty is distinguished by the type of mental action that is its paradigmatic expression, namely, active imagining. Active imagining as a type of mental action is to be understood in terms of its being a distinctive kind of means to the performance of a variety of further learned behaviors identified as imaginative by philosophical and empirical theorizing. In particular, active imagining consists in the goal-directed (potentiation of the) construction of a discursive representation wherein the agent selects the content of *that* representation in light of the goal for *that* instance of construction. Active imagining

skill for doing *A* is minimal but exists insofar as she can do *A* or to abstain from doing *A* in accordance with the relevant intention.

as a type of mental action is thus partly defined in terms of the wider actions in which it participates and partly in terms of its unique contribution to the performance of those wider actions. Likewise, what ties behaviors like pretense, engagement with fiction, prediction or explanation of others' behavior, reasoning about possibility and necessity, episodic hypothetical thought, and imagining for its own sake together as types of *imaginative* project is each behavior's taking as a means the agent's active imagining, i.e., her selective (potentiation of the) construction of the appropriate discursive representation.

An upshot of my account of active imagining is that all imagining is subject to the imposition of top-down and lateral constraints specific to the aim with which the agent actively imagines. Therefore, all imagining is constrained. For instance, imagining as part of pretense is subject to a top-down constraint that the agent imagine the goings-on of the pretense in a way that motivates play. Such imagining is thereby subject to certain lateral constraints, for instance, that its content update appropriately in response to the performance of pretense actions (Gendler 2003, 2006a,b,c). Another example: imagining as part of engaging with fiction is subject to a top-down constraint that the agent imagine the world of the fiction. Such imagining is thereby subject to the lateral constraint that its content update appropriately in response to what is implied by the author's presenting facts about the relevant fictional world (Walton 1990). In successfully taking up a game of pretense or a work of fiction, the agent determines that she actively imagines appropriate to playing the game or engaging with the fiction, respectively.

Recall Sant'Anna's claim that remembering differs metaphysically from imagining in part because an agent can inject whatever novel content she likes into imagining, whereas she cannot do so with remembering. This might seem to suggest that at least some kinds of imagining are unconstrained. Even if she cannot inject whatever content she likes into an instance of, say, predicting or explaining others' behavior and still count as doing that, still, at the limit, the agent appears to have complete latitude over what she imagines when she is imagining for its own sake. But even here imagining is subject to the imposition of both top-down and lateral constraints. In bare imagining for its own sake, the agent must imagine in line with her intention and she must have the conceptual repertoire or perceptual prototypes necessary to at least potentiate the selective construction of the intended discursive representation. The former is a top-down constraint and the latter is a lateral constraint. The former is imposed directly by the agent's intending, say, to visualize a dancing banana and the latter is imposed indirectly by this intention. She satisfies these constraints in her at least potentiating the relevant bit of construction through the exercise of her knowledge-how to construct an image in general, construct images of bananas, and construct images of something as dancing. That is, the agent determines that

she visualizes a dancing banana. The extent to which it is true that the agent can imagine whatever she likes is therefore coincident with the extent to which she can efficaciously imagine for its own sake. Outside of the context of imagining for its own sake, Sant'Anna is wrong to suggest that the agent can inject whatever novel content she wants and still count as, for instance, imagining being the person whose behavior she is trying to predict or explain. Moreover, imagining for its own sake should be understood as a limiting case that shows the boundaries of the freedom of imagination, not the paradigmatic exercise of the faculty. Even when the agent is imagining for its own sake, she cannot imagine just whatever she likes.

If it is right that all imagining is constrained then, *a fortiori*, epistemic imagining is constrained. What makes an instance of imagining epistemic is its serving as a means to the completion of an imaginative project that has as its aim the justification of some belief. So, epistemic imaginings are not distinguished by being subject to the imposition of constraints—lateral or otherwise—but by being subject to the imposition of specifically *epistemic constraints*. At least some philosophers of imagination working on the epistemic relevance of imagination seem to agree. For instance, Kind (2016) lays out two constraints specific to reasoning that makes use of imagination. The first requires that the agent hold fixed the relevant details about the actual world that she does not modify in the imagined scenario. Kind calls this the "reality constraint." The second requires that the agent imagine the causal or logical consequences of any modification that she does make. Kind calls this the "change constraint." Kind uses intelligent machines from John W. Campbell's 1932 science fiction short story "The Last Evolution" as examples to illustrate perfected forms of the imposition and satisfaction of these constraints. Their superhuman epistemic imagining allows these machines to hold fixed any fact whatsoever and to work out every minute causal and logical consequence of whatever modifications they make. In Campbell's short story, they use this power "to figure out not only how to counter a devastating attack from an alien force but also how to transcend their material existence (the titular "last evolution" occurs when the machines evolve into beings composed of pure energy)" (Kind, 2016: 149-150). Similarly, Myers (2021a) lays out a constraint specific to epistemically useful visualization. In particular, he holds that epistemically useful visualizations are justified justifiers. Recall the sofa visualization case. According to Myers, for the agent's visualization to justify her belief that the sofa will (not) fit through the door frame, her imagery must be constrained by her justified perceptual beliefs about the spatial dimensions of her sofa and door frame as well as by relevant justified background nomological beliefs.

Both Kind and Myers are here setting out lateral constraints on epistemically useful imaginings. Neither the reality constraint, change constraint, nor the constraint that epistemically useful visualizations stem

from justified perceptual beliefs is imposed in the first instance by the agent. Nonetheless, they are all indirectly imposed by the agent's intention to figure out, say, whether her sofa will fit through her door frame, how to counter a devastating attack from an alien force, how to transcend her physical form or whether such a transformation is even metaphysically possible, and so on. It is the agent's taking up these kinds of epistemic projects that ultimately imposes the relevant lateral constraints on her constructive activity and it is her at least potentiating the selective construction of the corresponding veridical or otherwise justified discursive representations that satisfies those constraints. The agent thus determines that she figures out, say, that her sofa will (not) fit through her door frame, how to fend off the aliens, how to transcend this mortal coil, etc. Her figuring these things out are in all likelihood non-basic actions of hers and her determining that she does them depends on her determining that she imagine, e.g., her sofa and door frame with their actual spatial dimensions, their hypothetical interaction, and the relevant nomological regularities.

I want to suggest that epistemically useful episodic thought in particular is unified by being underwritten by suitable happenings of exploring or navigating local regions of modal or temporal space through (potentiated) scene construction. Episodic hypothetical thought is (typically) underwritten by navigation to or exploration of sets of relatively nearby possible worlds. In the sofa visualization case, the set of worlds that the agent locates in her act of visualizing is extremely like the actual world and, so, very close to it. In that case, both the agent's navigating to or exploring the relevant set of worlds and her way of navigating to exploring it, namely, by self-imposing and subsequently satisfying the relevant constraints through her constructive activity, are informative. She does not have to "go far" to figure out whether her sofa will fit and she can use the sofa and door frame to scaffold her visualization, since they are both present to her. Both contextual facts speak to the plausible accuracy of her resulting visualization, should she decide to make use of the relevant scaffolding and to keep her imagining to the most nearby possibilities. Likewise, the relative closeness of such worlds is informative. Indeed, in the case under consideration, it is partly in recognition that the possibility that the sofa fits through the door frame is more likely than the possibility that it does not that the agent rightly infers or recognizes that she can fit the sofa through the door frame.

If this is the right way to think both about episodic thought as a subset of epistemically useful imaginative projects as well as the imposition and satisfaction of constraints involved therewith, then at least intentional episodic remembering turns out to be the agent's at least potentiating the construction of scenes underwritten by navigating or exploring her actual personal past. Recall the specific-past-experience constraint. It is the lateral constraint on episodic remembering that the content involved in mnemic activity appropriately

match the corresponding specific past experience. It is likewise the constraint that accounts of memory as a mental action seem to have in mind in endorsing the constraint commitment. On the account of active imagining I develop, episodic remembering is the agent's at least potentiating the construction of a discursive representation and her selection of its content as a means of re-experiencing an event underwritten by navigating or exploring her actual personal past. It is in aiming to re-experience the event that she ultimately self-imposes the specific-past-experience constraint. And it is by her successful (potentiation of the) construction of a discursive representation of the appropriate scene(s) sufficient to activate corresponding perceptual recognitional skills and underwritten by her navigation to or exploration of that event that she satisfies that constraint. It might well be that memory traces as well as the well-functioning of first-order mnemic mechanisms and metamemory function together to facilitate such navigation or exploration in part by imposing the specific-past-experience constraint. If so, the imposition of the constraint is directly analogous to the imposition of the relevant lateral constraints on visualization and on arm raising. In all three cases, although it is not the agent who imposes the lateral constraint in the first instance, she is the one who ultimately imposes it by virtue of her aiming to raise her arm, visualize, or remember, respectively. Moreover, it is she who satisfies that relevant lateral constraint by raising her arm, visualizing, or remembering.

This way of conceiving of the specific-past-experience constraint resolves the apparent tension between the constructive and constraint commitments within the Argument from the Constructive Nature of Imagination. Here is the argument again:

The Argument from the Constructive Nature of Imagination

P1) Imagining is active in that, other things being equal, the agent imagines when she intends to imagine and imagines what she intends to imagine.

P2) If remembering is the same kind of agential activity as imagining then, other things being equal, the agent remembers when she intends to remember and remembers what she intends to remember.

P3) Remembering is the same kind of agential activity as imagining.

C:) Other things being equal, the agent remembers when she intends to remember and remembers what she intends to remember.

In presenting my account of active imagining, I have given a sketch of why (P1) is true beyond accepting that the imagination is free. The agent is intending to do something using her imagination as a means and is very often successful at making appropriate use of her imagination (*cf.* Walton 1990: 20). She thus imagines when she intends to, say, pretend that the floor is lava. And she imagines what she intends to, say, *that* the floor is lava.

Let's turn to (P3). An obstacle to substantiating (P3) was the apparent tension between the constructive and constraint commitments. That tension arose because the imposition and satisfaction of the specific-past-experience constraint seemed to be outside of the agent's control. By contrast, I have been arguing that it is the agent who ultimately imposes the specific-past-experience constraint and that it is through her appropriate (potentiated) scene construction that she satisfies the constraint. That is, I have been suggesting that the agent determines that she remembers. On this suggestion, the imposition of the constraint is ultimately due to her undertaking an act of remembering. After all, she can stop imposing that constraint, say, by abandoning her intention to remember. Moreover, she can modulate the constraint both in specific instances and over time. For instance, she can respecify her intention from remembering her last birthday party to remembering it in reverse order, remembering what the cake was like, remembering what it was like to be surrounded by her friends, and so on. She can likewise repeatedly recall the event, making it easier to recall later and effectively making the act of remembering her birthday less stringently constrained. The agent's abilities to release herself from the specific-past-experience constraint and to modulate it suggest that it is she who ultimately imposes it on her own constructive activity. In effect, the agency exhibited in remembering that Sant'Anna (2023) wants to preserve is sufficient to claim that it is both constrained and constructive when considered from my account of active imagining.

If all of this is right, then it is the case both that episodic memory is constrained by the past and that it is appropriately like imagining in virtue of being constructive. Memory's being constrained by the agent's actual personal past is a function of its being the particular kind of epistemically useful imaginative project that it is. Nonetheless, it is by engaging in the very same activity that the agent imagines fantastical scenarios, plays pretend, and remembers. Funneling that activity into the (potentiated) construction of the appropriate scenes allows the agent to re-experience the relevant event and to thereby draw from that construction justified beliefs about that past. The imposition of such epistemic constraints—specifically the specific-past-experience constraint—is fully consistent with the fact that the agent's at least potentiating the selective construction of discursive content satisfies those constraints because it is ultimately the agent who imposes them. Thus, we substantiate (P3) of the Argument from the Constructive Nature of Imagination by assimilating episodic remembering as a kind of epistemically useful imaginative project on a par with, e.g., episodic hypothetical thought. From this we can infer (C).

To close out this section, let us reconsider the accounts of memory as a mental action and see whether they can adopt my suggestion for substantiating (P3). Starting with Hopkins, recall that he holds that the

content of memory is causally controlled by the corresponding past event. It might seem that he cannot adopt my suggestion, since to be causally controlled by the past event seems to exclude the possibility that it is the agent's selecting the content of the relevant discursive representation that fixes the content of her memory. Moreover, Hopkins has recently argued that the capacity to imagine and the capacity to remember are distinguished by the fact that the former is a power to represent in a general way whereas the latter is a power to represent particulars (2023: 220-222). Nonetheless, I do not take either of these claims to exclude Hopkins from taking my suggested substantiation of (P3) and more fully endorsing metaphysical process continuism. After all, he can cash out causal control in terms of indirect self-imposition of the specific-past-experience constraint. Moreover, he elsewhere (2023, 2024) argues that *everything* about the imagery formed in episodic thought is determinable by what the agent is committed to in imagining thus and that the capacity to imagine envelopes the capacity to remember. Taken together, Hopkins could well hold that the agent self-imposes the specific-past-experience constraint in intending to remember and commits thereby to at least potentiating the construction of the appropriate scene, deploying her capacity to imagine a scene like the one originally experienced *in* deploying her capacity to remember that experience where her doing so is necessarily basic.

Next, recall that Arango-Muñoz and Bermúdez hold that the agent indirectly controls what she remembers through her practical responses to the metacognitive feelings that she has in light of the unfolding of first-order scene construction. It might seem that they cannot adopt my suggestion, since the role that they attribute to metacognitive feelings seems to be the same that Sant'Anna attributes to source-monitoring. That is, it might be that Arango-Muñoz and Bermúdez would resist my suggestion on the grounds that the metacognitive feelings specific to remembering generate an asymmetry between remembering and imagining of the kind Sant'Anna takes to distinguish them as processes. Granted. That they could in principle align themselves with Sant'Anna does not show that Arango-Muñoz and Bermúdez are barred from taking up my suggested substantiation of (P3). After all, they do not say exactly in virtue of what the metacognitive feelings specific to remembering are activated or what directs them to the particular memory the agent is aiming to retrieve. In which case, they could take my suggestion onboard as filling out some of the background details of their account of memory as a mental action. Moreover, given that they explicitly appeal to the control we are supposed to enjoy over imagining as a point in favor of the claim that we enjoy a similar kind of control over remembering, they should probably avoid aligning themselves with Sant'Anna and should take up my suggested substantiation instead.

Finally, recall that, on my account, the agent can but often does not directly intervene on her mnemic processing. Not only is my account well suited to substantiating (P3) of the Argument from the Constructive Nature of Imagination by assimilating episodic remembering as a kind of epistemically useful imaginative project. A non-reductive anti-causalist version of my account of memory as a mental action drops out of my account of active imagining! Recall that on my account of active imagining, to imagine is to at least potentiate construction of some discursive representation and selection of its content as a means to performing some wider *learned* behavior like pretense, engaging with fiction, reasoning about possibility and necessity, episodic hypothetical thought, and so on. I contend that there are no inbuilt modules for the imaginative projects that an agent's imagining facilitates. Rather, these wider behaviors are skills that the agent acquires through practice-based learning and mastering of principles, methods, techniques, heuristics, and so on specific to each type of imaginative project. At least some of these principles, methods, techniques, heuristics, and so on themselves constitute lateral constraints on the agent's active imagining. Her practice imagining as part of embarking on the relevant type of imaginative project quite literally shapes her capacity to imagine over time such that its exercise eventually comes to automatically impose and satisfy the relevant lateral constraints. If, as I have suggested, remembering is an epistemically useful type of imaginative project then it follows on my account of active imagining that remembering is skillful. This result dovetails with a non-reductive anti-causalist version of my (2022) account of memory as a mental action.

I say that a non-reductive anti-causalist version of my account falls out of my account of imagining because the latter is also non-reductive and anti-causalist in light of being backed by the powers framework. In my (2022), I assume a generic form of causalism and try to defend the claim that episodic remembering is skillful even within this generic causalist account of action. My departure from causalism stemmed, in part, from my being convinced that at least some habitual (mental) actions did not require an attendant intention to (non-deviantly) control or guide them. At least, my conviction was that if an intention is understood as a truth-conditional propositional directive attitude then no such thing is required for the successful performance of at least some habitual actions. Recalling chapter 2, I can now give another reason why I ended up abandoning causalism about action. In particular, I realized that the control that Sant'Anna (2023) claims would have to be required for remembering to count as metaphysically continuous with imagining could not be had on any form of causalism. In Section 2.4.1, I put this in terms of how an agent relates to her action as agent. I said that the claim that intentional actions are constituted by (non-deviant) causal relations between a truth-conditional propositional directive attitude and the relevant corresponding happenings in the agent's neurocognitive

architecture, body, or the world outside of her body cannot make sense of how an agent knows her own actions. And I said that looking for a way to remedy this omission through notions of control or guidance were doomed to failure. For control in general on such views is just more (efficient) causation forming a nexus of loops. On such views, the activity of control mechanisms that gives rise to, say, the experience of thirst is not agential because its mechanisms are not identifiable with the agent. By contrast, the activity of control mechanisms in the so-called "executive system" that give rise to plans and the like is agential because its mechanisms are identifiable with the agent. Because I have been entertaining notions of control in this chapter in discussing Langland-Hassan's and Sant'Anna's arguments, I will give a defense of my shift to a non-reductive anti-causalist approach to memory as a mental action that deals more specifically with such notions.

The experience of thirst is supposed to come about by virtue of the functioning of two mechanisms in the subfornical organ that respond to levels of angiotensin II and sodium in blood plasma, respectively (Matsuda et al. 2020). As the levels of angiotensin II and sodium change, so too does the activity of these mechanisms. When levels of angiotensin II are high, "water neurons" in the subfornical organ are activated, leading to the experience of thirst. When sodium levels are low, these neurons are inhibited by GABAergic interneurons. The constant ebb and flow of these levels means that the firing rate and strength of inhibition of water neurons is constantly changing and that there is some threshold of activity above which the thirst response is supposed to be generated. That is, we have a causal loop among the activity of neural mechanisms involved in thirst, driving or suppressing such activity in response to levels of angiotensin II and sodium in the blood. Such mechanisms control thirst through the causal loop. But such control is not the agent's control. By contrast, neuronal activity in the executive system is supposed to be constitutive of agential control because this system supposedly "flexibly integrates and coordinates the workings of other subsystems *in a characteristically agential way*" (Hendrickx 2023: 3134; emphasis added). Yet, other control systems also "integrate information to generate flexible control over" such things as thirst, breathing, blinking, and so on (3127). So, the question is what makes the difference between integration and coordination that is agential in character and integration (and coordination) that is not.

A causalist at this point might well say that the difference comes from the fact that the former kind of integration and coordination is "top-down" or "endogenous" or realizes the agent's truth-conditional propositional directive attitudes such that those attitudes (non-deviantly) cause or bias the activity of what is thereby controlled by the executive system. But "top-down" and "endogenous" just mean that the activity constitutive of control originates in certain parts of the brain and not others. And pointing this out cannot help

make the distinction: why should the activity's originating in frontal and certain central areas of the brain make a difference? After all, the subfornical organ is in the center of the brain too. Moreover, it is not as though activity originating in the executive system is not itself responsive to what else is going on in the brain (or in the world). It is activated in response to other activity in the neurocognitive architecture that is somehow responsive to situations requiring planning, strategizing, maintaining attention, and so on. And such activity is very often mediated by perception. So, there are "bottom-up" or "exogenous" drivers of executive function just as there are such drivers of subfornical organ function. The choice of what counts as part of the relevant causal loops and what does not, then, does not stem from a clear line of demarcation around where neuronal activity is coming from. Finally, to say that what makes the difference is that activity in the executive system realizes truth-conditional propositional directive attitudes such that those attitudes (non-deviantly) cause or bias the activity of what is thereby controlled begs the question in favor of causalism. Sure, some cognitive neuroscientists will speak of "goal-states" and the like in describing what the activity of the executive system is supposed to realize. And some philosophers have claimed that such representational posits are useful or practically indispensable in neuroscience (for instance, Shea 2018). Yet, others disagree (for instance, Favela and Machery 2023). So, at least at this stage, causalists seeking to find refuge in cognitive neuroscience will not find anything definitive to turn to. Overall, then, appealing to "agential control"—where such control amounts to specific kinds of causal loops among the activity of certain mechanisms in certain parts of the brain—does not suffice to distinguish the kinds of causes that count as constitutive of intentional action from those that do not. And I doubt that it will in the future.

Returning to memory, the kind of control Sant'Anna (and others) have thought necessary to count remembering as metaphysically continuous with imagining is the same kind of control that some causalists have thought picks out (non-deviant) effects constitutive of intentional action. The agent cannot exercise such control over her mnemic content because of the activity of source-monitoring in memory. To have such control would require activity in the executive system to bias activity in first-order constructive mechanisms that conflicts with activity in the relevant source-monitoring mechanisms. And that conflict cannot be resolved except by abandoning the intention to remember or by avoiding injecting novel content into whatever mnemic content has already been formed. It cannot be resolved by the injection of such novel content and "ignoring" the activity of source-monitoring mechanisms on pain of lapsing into irrationality. But if the notion of agential control—understood as activity in the executive system—is not a stable foundation on which to build one's argument for metaphysical (dis)continuism then any such arguments inherit the shakiness of that foundation. I

have suggested that the notion of agential control does not provide a stable foundation because it does not suffice to distinguish what the causalist needs it to, namely, causes that (non-deviantly) make their effects constitutive of intentional actions and causes that do not. An agent's failing to have such control over the activity of her first-order constructive mechanisms, then, is no bar on claiming that the agent selects the content of her memories.[121] For such selection does not amount to that control.[122]

Instead, as I have argued with Springle elsewhere (under review), on the powers framework, selection amounts to at least potentiation of the relevant discursive solution systems for discursively representing certain contents. An agent's potentiating these discursive solution systems amounts to her having an embedding *de agendo* representation of an opportunity to construct a discursive representation whose use would in turn potentiate perceptual recognitional skills tied to the specific event in question in light of a need to re-experience that event. These perceptual recognitional skills are tied to that event in virtue of being formed at encoding and stabilized in part through consolidation. The causal history of the formation and stabilization of such capacities are ultimately what allow navigation to the relevant sliver of modal and temporal space. The need establishes the appropriateness of the potentiation of these recognitional skills. Such appropriateness in turn establishes the appropriateness of the content she selects in at least potentiating the formation of the relevant discursive representation. And she selects that content in light of the relevant need because she does not get it directly from perception or as a result of inference. Because I am not aiming to provide a full account of episodic memory in this chapter, I will leave this summation of Springle's and my view as a bit of gesturing towards a viable alternative to Sant'Anna, Langland-Hassan, and many other empirically informed philosophers of memory. The chapter will be successful if it provides a viable way of resolving the apparent tension between the constructive and constraint commitments. I have suggested that the powers framework gives us a way of resolving the tension by considering constraints in intentional action in general, active imagining, and epistemically useful imagining.

4.6 Conclusion

In this chapter, I have attempted to defend the claim that episodic remembering is a mental action. I drew out two commitments shared by accounts of memory as a mental action, namely, that episodic remembering is

[121] At least, it does not at this point block making this claim, given the current state of cognitive neuroscience.
[122] My mistake in my (2022) was to attempt to make space for agential control over what one remembers while acknowledging that the agent does not literally pick what is to be represented with complete latitude.

constrained by the personal past and that it is like imagining in virtue of being constructive. At first, these commitments seemed to be in tension with each other. How can remembering be constructive in nature if it is subject to constraints imposed by things outside of the agent's control? How can the agent's past control her remembering if she can selectively construct scenes as part of her mnemic activity? I resolved this tension by showing that the constraint specific to memory, namely, the constraint that mnemic content appropriately match the corresponding specific past experience, is at once self-imposed and satisfied by the agent's intentional remembering. Her determining that she remembers, say, her last birthday party just is this self-constrained activity. I argued that this activity is imaginative by virtue of being the (potentiation of the) selective construction of the appropriate discursive representation as a means sufficient to activate corresponding perceptual recognitional skills. All imagining is constrained by the aim with which the agent at least potentiates(/generates an intentional action acorn for) the selective construction of the relevant discursive representation. Epistemically useful imaginative projects have specifically epistemic aims, where the subset of such projects comprising episodic thought at least involves as suitable happenings the navigation of modal or temporal space. And episodic remembering is an imaginative project with the aim of re-experiencing an event whose completion is underwritten by navigating the actual personal past. Episodic remembering is thus a highly constrained, skillful epistemic imaginative activity.

5.0 Conclusion

Drawing from the broadly Rylean and Anscombean metaphysics of agency provided by the powers framework in chapter 1, I sought in chapter 2 to set out and defend an account of mental action that radically differs from those that are standard in the literature. According to that account, mental actions are the expression of mental skills and mental habits. Mental skills are rational dispositions consisting in networks of internally calculatively structured intentional action concepts for the construction and use of discursive representations. Mental habits are volitional dispositions that incorporate corresponding mental skills, and at least some are subsumed by or integrated with other mental skills. Like all human skills and habits, such dispositions are acquired and developed through rational transmission as part of the life cycle of the relevant mental skills and as part of an agent's inculcation into the relevant discursive social praxes. And like all human skills, mental skills are a teleological unity that develops in characteristic stages of *a priori* moments of doing, practicing, and teaching. The *a priori* moment of doing in turn consists in a teleologically unified process that develops in characteristic stages of *a priori* moments of activation of the relevant parts of the network of intentional action concepts composing the relevant mental skill, ongoing performance of the mental action, and completion of that action. These moments are backed by the agent's cogent practical cognizing or, what amounts to the same thing, her ordering the means and ends of the relevant action in a calculation what to do that is backed by what I call "practical want." I identify practical cognizing with the intention with which the agent acts. And practical cognizing, when cogent, is constitutive of exercises of practical knowledge. And the exercise of practical knowledge is, in turn, constitutive of the intentional action that the agent performs.

What distinguishes mental actions from non-mental actions is the most general type of problem that each functions to solve and the most general type of means required to solve that problem. Mental actions ultimately function to solve the problem of theoretical and practical rational self-integration of an agent by means of the construction and use of discursive representations. Mental actions "proximally" function to provide mediated practical access to targets of action by means of the construction and use of such representations. When mental actions are overt, they are the actual construction and use of such representations sufficient to activate corresponding recognitional skills. When they are covert, mental actions are (merely) the potentiation of such construction and use sufficient to activate corresponding recognitional skills. Finally, such potentiation is the tokening of embedding *de agendo* representations (or intentional action acorns) of discursive representational vehicles as constructable or useable such that those vehicles stand in or go proxy for the target that the agent thereby discursively represents. Discursive representations are, in general, embedded by

corresponding embedding *de agendo* representations and their veridicality conditions are thus enveloped by their appropriateness conditions. The activation of recognitional skills as part of mental actions is the tokening of a multitude of interrelated embedding *de agendo* representations that function to discursively represent the same content.

In chapter 3, I drilled down on imagination, providing an account of it that derives from the powers framework's account of mental action. According to the Skillful Action Account of Imagining, to imagine is to at least potentiate the construction of a discursive representation and the selection by the agent of that representation's content as a means to the completion of learned behaviors that, following the literature, I labeled "imaginative projects." Imagining is unified by being the skillful selective constructive means to the completion of such projects and such projects are unified by taking such selective constructive activity as means. On the Skillful Action Account, the intentionality of imagining consists in the accordance of agent's selective constructive activity with her cogent practical cognition concerning how she goes about completing the relevant imaginative project. I argued that the account can handle seemingly passive instances of imagining like automatic imagining, involuntary imagining, and mind-wandering either by understanding them as in fact active or by treating them as noncentral or derivative.

In chapter 4, I drilled down further on episodic memory, providing a way for accounts of memory as a mental action to resolve the apparent tension between holding that memory is both constructive in nature and constrained to representing particular events from the personal past. I argued that the key to resolving the tension was to recognize that all rational agency is constrained by the agent herself. Recharacterizing cogent practical cognizing as the self-imposition and satisfaction of constraints on action, which I labeled "self-determination," I argued that imagining in general is self-determined as is epistemically useful imagining. In which case, there is no tension between the imposition of the constraint to construct a scene of the particular event that the agent aims to recall and her selection of the content in (the at least potentiation of) the construction of that scene. The constraint is self-imposed by the agent's forming the intention to remember the event and satisfied by her mnemic activity in light of that intention. I then argued that each of the three accounts of memory as a mental action that I consider can adopt my proposed way of resolving the apparent tension.

There is a great deal that I could not do in the dissertation. I realized fairly late in the process of writing it that a full defense of the powers framework requires a coherent account of rational transformation. That is, the framework needs an explanation of how inculcation into discursive social praxes and the rational

transmission of skill fundamentally changes human agents' perception, knowledge, and action. This need is all the more pressing given that I have claimed that the ground floor of the human mind consists in the initial mental skills and mental habits that an agent acquires in growing into the age of reason or, what amounts to the same thing, in coming into the space of reasons. Importantly, such an account of rational transformation should also make clear how human mindedness is nonetheless continuous with nonhuman animal mindedness without allowing that rational capacities are merely additional faculties that sit atop non-rational capacities. Giving such an account is a dissertation or book unto itself, if not a life-long project (Eric Marcus, for instance, has published several papers trying to unpack rational transformation and, to my knowledge, is still deeply engaged in this work). If given the chance, I hope to contribute to the burgeoning literature on rational transformation.

Similarly, I was unable to provide substantial critiques of causalism about action in general or causalist accounts of mental action in particular (however, see Appendix F). What critiques I did give in the main text were meant only to make room for the powers framework account of mental action to stand as a viable account of mental action. Part of showing that it is a viable account took the form of arguing that it is in agreement with the Neutral Account of Mental Action. So, the target of the powers framework of mental action is the same as that of its causalist rivals. The other part was putting pressure on those rivals. But, again, I do not take myself to have provided any knockdown arguments against causalism in general or causalist accounts of mental action in particular. If given the chance, I'd like to provide more thoroughgoing critiques of causalist accounts of mental action. I was also unable to defend the powers framework's definitions of various of our mental faculties. In chapter 2, I laid out its preferred method for type-individuating mental actions but did not defend that method or the accounts of deliberation, inference, imagination, memory, human perception, attention, and so on. Another lifelong project is to defend the way the powers framework type-individuates these mental actions. A third lifelong project is to defend the accounts of the different mental faculties that come out of such type-individuation. If given the chance, I'd like to work on both. I believe that these three projects can be tackled together in articles and books on deliberation, inference, memory, human perception, attention, and so on. I have started doing some of this work specifically with respect to episodic memory with Alison Springle (under review).

I was likewise unable to provide a full defense of the claim that the powers framework, despite requiring a theory of rational transformation, is nonetheless amenable to empirical inquiry into the mind. The framework would have it that such inquiry, properly conducted, seeks to understand the causal mechanisms

that underwrite our mental skills and mental habits. In contrast to much contemporary empirically informed philosophy of mind, it denies that such inquiry is in the business of helping us understand the metaphysics of mind. Of course, it agrees that empirical theorizing on some topic should to some degree constrain metaphysical theorizing on that topic. Where it departs from other empirically informed accounts is the degree to which philosophers of mind should be deferential to the mind sciences. After all, those sciences take at least some of their assumptions from the metaphysical theorizing that has already occurred. Critical engagement with such sciences does not require complete deference nor offhand rebuke of the relevant findings. Rather, such engagement requires careful treatment of such findings and a willingness to push back on underlying assumptions that might not be necessary or which might obfuscate what such findings really imply or entail. It is becoming increasingly difficult to do one's due diligence here: there are powerful incentives in both philosophy and science to over promise and to publish like there is no tomorrow. Such incentives militate against the slow, careful process of sifting through the empirical literature, talking with scientists about their work, and doing a proper sociology of science with a view to getting clear on what the empirical findings really show. Of course, such tasks are not to be carried out by any single individual or by individuals working independently of each other. Like science itself, critical engagement as part of a proper philosophy of science is a communal endeavor. But the point still stands that task is daunting. If given the chance, I'd like to contribute what little I can to that endeavor.

Finally, the account of imagining that I give in chapter 3 is admittedly programmatic. It sketches an account and concretizes it somewhat by focusing on imagining as part of pretense and engagement with fiction. But there is much more work to be done. First, there is providing the Skillful Action Account's take on the other imaginative projects. This is likewise a lifelong project. Second, there is providing the account's take on various puzzles in the philosophical literature on imagination. How does the account understand the paradox of tragedy? How does it understand the epistemic import of thought experiment in science and philosophy? How does it understand negative existentials or the contents of fiction? Can it solve the problem of other minds and, if so, how? Third, there is providing a complete account of the intentionality of imagination. While the account can take cues from work already done on this issue (as well as work done on the various puzzles associated with imagination), it should have something more to say than the qualified form of intentionalism that I provide in Section 3.4. Fourth, in line with the previous loose thread, there is accounting for the various malfunctions of the mechanisms that underwrite the capacity to imagine. Some of this work is empirical in nature and should be carried out by the mind sciences. But, if given the chance, I'd like to contribute to interpreting results of that

empirical work and understanding it in light of the Skillful Action Account. Finally, there is fitting the powers framework in general, its account of mental action, and the Skillful Action Account of Imagining with logics of action. I do not take myself to have the training appropriate for completing such a task. But I hope to contribute what I can to its completion.

Given space constraints, there are a variety of issues that I could not touch on in the main text. I have written five Appendices to fill in details that I think are crucial to defending the powers framework but which I could not justify including in the first chapter. Appendix A gives some historical background on the notion of practical knowledge that Anscombe makes use of in various works. Appendix B provides a more thorough explication of what the powers framework takes to be the form of practical reason. Appendix C addresses the fallibility of practical cognition in human beings. Such fallibility consists in *prima facie* failures of such cognition or "errors in judgment" as well as mistakes in performance. It seeks to put into relief how robust and flexible practical cognition actually is by acknowledging that, while such cognition is not in every case perfectly cogent, nonetheless, it is not as fragile or brittle as many who have rejected the so-called "knowledge thesis" or "cognitive condition" have made it out to be. Appendix D furthers this aim by arguing that counterexamples of unwitting success like the carbon copier are, in fact, no threat to the knowledge-thesis or cognitive condition. Appendix E argues that counterexamples to the so-called "Simple View" are likewise, in fact, no threat to the powers framework, which adopts a sophisticated form of the Simple View. Finally, Appendix F considers accounts of mental action from Levy (2023), Wu (2023a,b), and Proust (2009, 2013a,b) and provides more substantive criticisms of them than I could fit into the main text. If the reader found that chapter 1 was lacking in regards to its explication and defense of Anscombean practical knowledge, I highly recommend reading the Appendices A-E. If the reader found that chapters 2 or 4 were lacking in regards to their treatment of reductive causalist accounts of mental action, I highly recommend reading Appendix F.

Appendix A: A Brief(ish) Historical Overview of Practical Reason and Practical Knowledge

Schwenkler (2015, 2019) has argued convincingly that we should understand Anscombe's account of practical knowledge as drawing heavily from both Aristotle and Aquinas, the latter himself drawing heavily from the former.[123] Given that the powers framework adopts a distinctively Anscombean account of intentional action, it too ends up indebted to Aristotle and Aquinas. From Aristotle, Anscombe takes the structure of practical reasoning or "'practical syllogism' which means the same thing" (1957/2000: §33). I claimed in chapter 1 that we should not identify practical reason as Anscombe understands it with the mental act of deliberating. Part of the reason for this is that, at least in cases of action involving perception, the "form of a calculation what to do" is already exhibited in a rational agent's manifesting her rational and volitional dispositions (§33). By definition, such dispositions are already internally calculatively structured and are activated on the relevant occasions by there being some target that the agent is in perceptual contact with and some need whose satisfaction would involve the appropriate type of interaction with that target. The calculative order so activated is further attuned to the situation at hand in light of the unfolding of the action once it has developed from activation to ongoing performance.

The point for now is that, recalling chapter 1, practical reason is very different in its content, form, and function from theoretical reason (Small 2022). The objects of practical reason are future contingents affectable by the agent, whereas the objects of theoretical reason are anything that could be true or false and which the agent could consider. The form of practical reason is that of calculating to an immediate action that has its point in getting the agent what she wants, whereas the form of theoretical reason is that of considerations of evidence or reasons for or against the truth of some conclusion. The function of practical reason is settling on action that closes the spatial, temporal, or conceptual distance between the agent and what she wants, whereas the function of theoretical reason is settling on the answer to some question, the production of knowledge of something as being the case. Given its objects, form, and function, the minimal standard of correctness for practical reason is goodness preservation. By contrast, given its objects, form, and function, the minimal standard of correctness for theoretical reason is truth preservation.

---

[123] Anscombe's daughter, Mary Geach, has attested to the extent of her mother's indebtedness to Aquinas, writing "Anscombe drew on his thought to an unknowable extent: she said to me that it aroused prejudice in people to tell them that a thought came from him: to my sister she said that to ascribe a thought to him made people boringly ignore the philosophical interest of it, whether they were for Aquinas or against him" (2005: xix). To give Aquinas's account of practical knowledge pride of place in *Intention* (§48) rather than bury it completely, then, suggests that Anscombe was likely going out of her way to signal her indebtedness to him concerning the notion.

These contrasts between practical reason and theoretical reason that Anscombe is drawing from Aristotle brings us to what she takes from Aquinas. While what she cites of Aquinas in *Intention* is minimal, as Schwenkler has argued, Anscombe's debt to him is substantial (2015, 2019, especially chapter 6). In particular, Aquinas draws three features of practical knowledge from Aristotle's discussion of practical intelligence in *De Anima* that Anscombe implicitly makes use of in *Intention*. I quote Schwenkler (2015: 11-12; original emphasis) at length:[124]

> [1. (**Object**):] First, as Aristotle claims that the object of appetite is always something attainable, or a "good that can be brought into being by action" (433a29–30), so Aquinas writes that the *objects* of practical knowledge must be things "producible by the knower": thus, *e.g.*, human knowledge of God or the natural world is necessarily [speculative], whereas it is possible for us to have practical knowledge of artifacts.
>
> [2. (**Form**):] Second, following Aristotle's claim that practical thought "calculates means to an end" (433a14), Aquinas identifies a practical *mode* of knowing, in which producible things are considered "as producible." By way of contrast Aquinas gives the example of an architect who "defines, analyses and examines the qualities proper to houses in general"—in this case the *object* of his knowledge is something that lies within his power, but the *mode* in which he knows it is not a practical one, as it involves only ascertaining what a house is, without any concern with how to make one.
>
> [3. (**Function**):] Finally, following Aristotle's remark that practical thought "differs from [speculative] thought in the character of its end" (433a15–16), Aquinas says that knowledge can be either [speculative] or practical in terms of its *end or purpose*. According to Aquinas, knowledge is practical in this respect insofar as it aims at production or some other form of action, and [speculative] in this respect insofar as it aims just at "the consideration of truth": thus, *e.g.,* the knowledge of builders who "consider how some house could be built, not with a view to building it but merely for the sake of knowing" will have a [speculative] end or purpose even though it is practical in its object and mode—such knowledge is what we call "idle speculation" about how something might be done, because it is not, as Anscombe puts it, "on active duty" (*I*, 60). By contrast, for knowledge to have a practical end or purpose is for the knowledge to have the aim of getting something done.

Practical knowledge for Aquinas, like practical reasoning for Aristotle, differs from speculative knowledge with regard to its objects, form, and function. What Aquinas adds is a set of divisions between kinds of knowledge that follow from these three distinguishing features. Thus, he writes "some knowledge is [speculative] only; some is practical only; and some is partly [speculative] and partly practical" (this and the following quotations are all from *ST* I, q. 14, a. 16c). He goes on to say that knowledge is "merely [speculative]" when its object is something "not operable by the knower; such is the knowledge of man about natural or divine thing[s]." Knowledge is "partly [speculative] and partly practical" when its object is something operable by the agent but "is [speculative] either in its mode[(/form)] or as to its end[(/function)]." Knowledge is "simply practical" when

[124] See also Schwenkler (2019: 158-159).

its object is something operable by the agent and "when it is ordained to an operative end," that is, when its form and function are such as to produce the agent's end by virtue of her taking what she understands as a means to it, namely, her intentional action.

We can divide the territory thus, from knowledge that is most speculative/least practical to knowledge that is least speculative/most practical:

> **Purely Speculative Knowledge:**
> Knowledge of things that do not lie in one's power based in considerations of fact which reveal the nature of those things.
>
> **Practical Content, Speculative Form & Function:**[125]
> Knowledge of things that lie within one's power based in considerations of fact which reveal the nature of those things.
>
> **Practical Content & Form, Speculative Function:**
> Knowledge of things that lie within one's power based in considerations of how one could bring those things about for the sake of revealing how one could bring them about.
>
> **Practical Content & Function, Speculative Form:**
> Knowledge of things that lie within one's power based in considerations of fact which reveal the nature of those things for the sake of later bringing them about.
>
> **Purely Practical Knowledge:**
> Knowledge of things that lie within one's power based in considerations of how one can bring those things about for the sake of bringing them about.

Aquinas's distinctions between speculative knowledge, practical knowledge, and mixed kinds of knowledge are important here because Anscombe's aim between §§4-49 (≈ 84 pages of a 94-page book) is accounting for intentional action and the intention with which a person acts and because she takes it that what accounts for both is what I have here labeled "Purely Practical Knowledge." When and only when an agent has purely practical knowledge of her doing *A* and of the point of her doing *A* does her doing *A*-ing count as her intentional action.[126]

---

[125] Neither Aquinas nor Schwenkler explicitly identify this sort of mixed knowledge, since, I take it, Aquinas at least would still classify it as merely speculative. Schwenkler appears to make a nod to it in the cited text under Form in discussing the architect who merely considers the nature of a house. The case is ambiguous between this sort of mixed knowledge and the sort wherein the knowledge has a practical content and function but speculative form. In any case, it is worth setting out this additional mixed kind of knowledge, since both Aquinas and Schwenkler define purely/merely speculative knowledge with respect to its object.

[126] That said, in her explication of practical reasoning, Anscombe seems to have something like knowledge with a practical object and form but with a speculative function in mind when pointing out that one could go through the following

Given the powers framework's commitment to a distinctively Anscombean account of intentional action, the framework is committed to the dependence of intentional action on purely practical knowledge. That is, an intentional action is nothing other than the agent's knowingly taking some practicable means to her end. Yet, for this to be what she is doing, she must understand *what* she is doing under the concept of means (and ends) and must understand *why* she is doing it in relation to whatever good she conceives the end as falling under. Such understanding as embodied in her ability to structure her own agency, ordering her rational and volitional dispositions in accordance with achieving the relevant end by taking the relevant means, is nothing other than the exercise of her purely practical knowledge. The framework thus adopts the factualist understanding of practical knowledge (Schwenkler 2019: 99).

Moreover, purely practical knowledge has the calculative form that constitutes the means-end structure of intentional action in virtue of being backed by the agent's practical reason. It is only because the agent has discursive solution systems for explicitly ordering means and ends in mental acts of deliberation that she can thereby know of her doing *A* that it is her taking some means which any agent situated as she is can take and which is both productive and constitutive of her doing *A*. That is, the possibility of *human* practical knowledge depends on the capacity for practical reasoning. Yet, not every instance of an agent's knowingly doing *A* is an instance in which she deliberates. The "premises" of the relevant practical syllogism are "on active duty" in cases of knowledgeable action in the absence of prior or concomitant deliberation when the relevant rational and volitional dispositions have been appropriately internally calculatively structured through training and practice and when those dispositions are potentiated in light of the agent's (perceptual) recognition of practicable means

---

inference without being subject to charges of inconsistency or other defect in rationality if they thereupon did not act (1957/2000: §33):

> P1) Vitamin X is good for all men over 60,
> P2) Pigs' tripes are full of vitamin X,
> P3) I am a man over 60,
> P4) Here's some pigs' tripes,
> C?) So, here's some food that is good for me/is a means of getting vitamin X, which is good for me.

Anscombe's recognition that action is not necessitated whenever one identifies practicable means sufficient to some end seems to depend on holding that the calculative form of reasoning does not of itself characterize practical knowledge. I take it that she has Aquinas's distinctions in mind here and I take it that she thinks that, in this case, the man's considering how someone like him might get some vitamin X has a speculative function: "if a builder should consider how a house can be made, not ordering this to the end of operation, but only to know (how to do it), this would be only a [speculative] considerations [sic] as regards the end, although it concerns an operable thing [and has the calculative form of reasoning]" (*ST* I, q. 14, a. 16c).

to some end of hers (Anscombe 1957/2000: §33). The relevant cluster of intentional action acorns are generated in light of the agent's recognition of the relevant target(s) and need(s).

In cases of knowledgeable action backed by prior or concomitant deliberation, the relevant dispositions are not already appropriately internally structured at least for the case at hand. The agent in such cases must structure her own agency, ordering the relevant parts of her solution systems into means appropriate to achieving her end. She can do this by explicitly discursively representing the relevant order in an overt mental act of deliberation, saying, writing out, mapping, diagramming, etc. the calculation of what to do. The construction and use of such explicit discursive representations potentiate the relevant part(s) of the agent's rational and volitional dispositions, priming her to instantiate the order she discursively represents. The pattern of activation at issue might well be novel for her, constitutive of a novel network of means and ends that might well result in the formation of a novel solution system (however short-lived it might be). Or, she might structure her own agency in a covert mental act of deliberation, potentiating the construction and use of the relevant discursive representations. Such potentiation has the same effect. And, in either case, she might end up searching for practicable means as part of her concomitant overt or covert deliberation should the need arise. She might say or think "I can do *B* by means of doing *A*" and then search out for a "*this* that I can do" constitutive of her being in a position to do *A* (vis-a-vis *this*). Finally, according to the powers framework, in neither case should the agent's practical reason or practical knowledge be understood in "cognitivist" terms. That is, the agent's calculation is not realized by a belief, at least if what is meant by "belief" is a truth-conditional propositional attitude that is realized by the activity of the agent's neurocognitive architecture. In cases of knowledgeable action in the absence of deliberation and in cases of knowledgeable action backed by covert deliberation, the only representations on the scene are *de agendo* representations. In cases of knowledgeable action backed by overt deliberation, the relevant truth-conditional representations are discursive representations that are actually constructed and used and, thus, in principle intersubjectively available.

Appendix B: The Practicality of Practical Inference

In this Appendix, I unpack more fully the twelve ways that Small elaborates the teleological conception of practical reason as drawn from Anscombe's (2005) "Practical Inference" and Ford's (2016) "On What Is in Front of Your Nose." Some of what I say here repeats what I said in chapter 1. The point is to provide a more complete picture of the form of practical reason that I sketch there. So, I urge the reader to forgive the repetition, especially if connecting the repeating parts to my further unpacking Small's elaborations helps them see the picture of practical reason he presents more clearly. (1) Following Anscombe, Small restricts the content of practical reasoning to "future contingents affectable by action" (267). A form of this restriction is found in Aristotle and Aquinas as well: all four philosophers identify that we can only sensibly calculate what to do about things that are within our power to affect and which we have not already affected in the relevant way.[127] (2) Small introduces a formal tool for connecting these future contingents together in a chain running from the agent's objective to the immediate action she performs. He identifies it as the "means-connective" and says that the notion of a means is "a *formal* concept of *practical* thinking: a practical conditional 'connective' that is a specific determination of [the material conditional]" (268; original emphasis). Means (or iterations of the means-connective) connect what he calls "doables: it is doing *A* that is a means to doing *B*" (268). Doables are just the now familiar intentional action concepts. I follow him in representing the means-connective with "→", thus, "*I can* do *B* by means of doing *A*" is represented as "doing *A* → doing *B*" (274; original emphasis).[128]

Now, neither the means-connective nor the notion of intentional action concepts suffice to account for what Small, following Anscombe, calls "practicability," namely, an intentional action concept's being something the agent "'can just do', can do 'straight away', or 'immediately make true'", that is, being immediately insatiable in action (2022: 268-269). Identifying and thereupon acting on practicable means is

[127] One might object here that such content restrictions disallow giving a formal characterization of practical reasoning, since identifying valid forms of reasoning ought to be done independently of the content reasoned about. But such an objection begs the question against the teleological view of practical reasoning. The point of theoretical reasoning is to move from truth to truth in a principled way, either discovering what was there already in a truth previously known (deductive) or making ampliative inferences to the truth of facts not contained within those previously known (inductive). Given this *telos*, the proper object of theoretical reason is any statement that can be true or false. The *telos* of theoretical reasoning is such that its domain is arguably total. The mistake in the objection is to fail to see that the lack of domain restriction on theoretical reasoning is a product of its *telos* rather than a constraint on reasoning as such.

[128] The astute reader will notice that I have switched the order from the How-direction typical of an agent deliberating what to do to the Why-direction that Anscombe more often discusses in *Intention*. I follow Small in doing this, do it for ease of exposition, and do it because the Why-direction and How-direction describe equivalent calculations. See Ford (2015).

what concludes practical reasoning. That is, practical reasoning ultimately takes on the following formal shape (adapted from 272):

Table B.1 The Form of Practical Reasoning
On the left-hand side, depicting the informal characterization of a calculation that doing *A* (vis-à-vis *this*) is a means to doing *D*. On the right-hand side, depicting the formal characterization of that calculation.

I can do *B* by means of doing *A*, and thus	P1. [doing *A* → doing *B*]
I can do *C* by means of doing *B*, and thus	P2. [doing *B* → doing *C*]
I can do *D* by means of doing *C*, and thus	P3. [doing *C* → doing *D*]
I can do *D* by means of (doing *C*, which I can do by means of doing *B*, which I can do by means of doing *A*, which I can do by means of) doing *A* vis-à-vis *this*.	P4. [doing *A* vis-à-vis *this* → doing *D*]

Whereupon the agent does *this*, i.e., does *A here and now*. And depending on whether *D* fits within a wider description of what the agent is now doing or not, the agent's so acting will thereby be either her *D*-ing or her *going to D* (2012: 166*ff.*), respectively.

(3) Following Anscombe, Small identifies practicable actions as those that are sufficiently specific relative to the agent and involve particulars as proximal targets (2022: 269). Ends can be general or particular and can differ in their degree of specificity relative to an agent. For instance, an agent can want to make *some* hut inhabitable (general and generic), to make *this* hut inhabitable (particular and generic), to move her belongings and things that she can use straightaway that suffice to keep her living into *some* hut (general and specific), or to do so with respect to *this* hut (particular and specific).

Table B.2 Generic and Specific Versus Generality and Particularity
Depicting the orthogonal distinctions of generic versus specific and general versus particular as they relate to action.

	General (categorical & absolute)	Particular (categorical & absolute)
Generic (graded & relative)	Make *some* hut inhabitable	Make *this* hut inhabitable
Specific (graded & relative)	Place belongings and usable survival tools into *some* hut	Place belongings and usable survival tools into *this* hut

(4) Following Anscombe and Ford (2016), Small holds that means move towards full particularity and sufficient specificity. If the agent is to make *some* hut inhabitable, she will need to find a particular hut and make *it* inhabitable (general to fully particular). Moreover, if she is to make *that* hut inhabitable, she must identify ways of making it inhabitable that are within her power as it stands, e.g., moving her belongings in, gathering

wood and kindling, using a striker to start a fire, and so on (generic to sufficiently specific). Assuming she can 'just do' these things as described 'straight away' or 'straight off' or can 'immediately make [them] true,' they are sufficiently specific means to making the relevant hut inhabitable (Anscombe 2005).

(5) Following Ford (2016), Small (2022) holds that perception plays an essential role in practical reasoning by playing an essential role in specifying and fully particularizing practicable means. Indeed "the acquisition of practicable knowledge of particular means requires a literal *search*, a physical intentional action that essentially involves sense-perception, the success of which requires the practical perceptual recognition of particulars" (270; original emphasis). According to Small, to reach a practicable means that she can enact, the agent must make perceptual contact with particulars in the environment whose arrangement is such that the agent recognizes an opportunity to implement *that* means and recognize it as such, i.e., as a "*this* that I can do." Recall that, in Springle's framework, the application of exploratory solution systems functions to generate further *de agendo* representations whose exploitation by the agent in turn modifies solution systems she already possesses or develops new ones altogether. And recall that I said that such modification and development is likely the rule rather than the exception. It is due to the essential role that perception plays in specifying and fully particularizing practicable means that modification or development of solution systems is the rule. That said, modification or development of sub-solutions might well be extremely minor and short-lived in cases where the agent is applying a well-worn solution system. These are small alterations to an otherwise fairly stable skill or habit. Even so, adopting Springle's framework and this claim about perception suggest that human agents (at least) are constantly making such small alterations. And this fits well with thinking of such dispositions as skills and intelligent habits.

(6) Likewise following Ford, Small holds that the premise identifying practicable means via perception—i.e. (P4) of the aforementioned syllogism—involves also recognizing "obstacles and opportunities" to the performance of the relevant action (272). Small adds that what counts as an obstacle or opportunity is relative to the agent's skill. As he puts it "ten-knot winds present an opportunity to an advanced sailor but an obstacle to a novice" (272). Returning to Springle's framework, what modifications or developments can or must occur for the agent to succeed in applying a solution system and the degree of modification or development allowed or required for that success, then, depend not just on the agent's identifying a "*this* that I can do." They depend as well on the shape *in concreto* of the skills or habits that she aims to manifest at the time of performance. The shape *in concreto* of the advanced sailor's disposition for sailing will allow for a variety of optional modifications of sub-solutions in light of ten-knot winds. Yet, he must elect at least one of those

modifications in manifesting his disposition. By contrast, the shape *in concreto* of the novice's disposition for sailing will likely require her to heavily modify her sub-solutions or develop entirely new ones in order to effectively manifest that disposition.

(7) Adding to Ford's (2016), Small argues that an agent's knowledge of any means-connection between intentional action concepts as practicable requires in addition that she "actually believe [that *she* can do *B* by means of doing *A*]: the possibility it represents as really practicable is one [her] taking of which requires deploying knowledge of this very fact" (272). He continues:

> If I am doing *B* by means of doing *A*, then I am doing *A* in order to do *B*, *on account of my insight that doing A is a way for me to do B*: if I do not know that I can do *B* by means of doing *A*, then it is not true that I *can*. (It may be true that I *could*, if only I *did* know!) (272-273; original emphasis).[129]

Small later identifies the type of fact that an agent in this position believes as the type involved in practical reasoning proper (275). Again, on the powers framework as I am elaborating it, such belief (or knowledge) should not be thought of as an additional occurrent mental state alongside the agent's internally calculatively structured habits and skills. Rather, such belief (or knowledge) just is the agent's possessing the relevant intentional action concepts and their being related by the means-connective. That is, an agent's possessing that belief just is her possessing the relevant internally calculatively structured habits and skills. Again, such beliefs are "true" whenever manifestation of that structure in action would suffice to close the distance between herself and what she practically wants were she to practically want it. Such beliefs are "false" whenever manifestation of that structure in action would not suffice to close that distance.

(8) Adding to Ford's (2016), Small argues that, in addition to obstacles and opportunities, the specificity sufficient for an action identified as practicable to actually be practicable depends not just on context (271) but on the agent's skill as well (272). In fact, as an agent becomes more skilled, she becomes more able to implement both more and less specific means to her ends. On the one hand, the master pianist can straightaway play *Wrong Note*, whereas a merely proficient pianist needs to identify more specific means in order to go about playing it. "Play *Wrong Note*" is practicable for the master, not for the proficient pianist. Another way to put the point is that the master pianist's rational disposition for playing has within its network of intentional action concepts one for playing *Wrong Note*. Her possession of such a concept is a determination of the shape of her

[129] This is an affirmation of the knowledge thesis, cognitive condition, or factualist understanding of practical knowledge. See chapter 1 and Appendices A, C, D, and E.

capacity to play *in concreto*. This concept, in turn, is constituted by a subsumed habit that incorporates the skill for playing *Wrong Note*. And this latter rational disposition is itself articulable into further intentional action concepts and means-end relations between them. Some of these concepts will be constituted by further subsumed habits and incorporated skills. What's important here (and below) is that the master pianist has a habit for playing *Wrong Note* as part of the determinate shape *in concreto* of her skill to play. By contrast, the proficient pianist will not have a subsumed habit to play *Wrong Note* despite having habits and skills whose exercise in the case under consideration is sufficient to count as his playing *Wrong Note* intentionally. He might well have a skill to play *Wrong Note*, assuming he has played it intentionally before. But he has yet to achieve the shape *in concreto* of his skill to play that the corresponding disposition of the master pianist possesses. Such modifications of expert skills are time consuming and hard-won.

On the other hand, the master pianist might well be able to play *Wrong Note* playfully or as though she were really getting the notes wrong, whereas a proficient pianist, if successful, is likely not going to be able to do anything more specific than play the piece. "Play *Wrong Note* playfully" is practicable for the master, whereas it is not something that the proficient pianist can even properly aim at. The same claims about the differences in the shapes *in concreto* of the master's and proficient pianist's rational dispositions to play apply here as did concerning the ability to play *Wrong Note* with one exception. Namely, the proficient pianist will not even have habits or skills whose exercise on the relevant occasion would suffice for him to be playing *Wrong Note* playfully.

(9) Following Anscombe, Small identifies three types of fact at issue for practical reasoning. These types are as follows: "type-(i)" facts are those of the form "*one* can do *A* by means of doing *B*", i.e., "doing *A* 'is' a means of doing *B*." "Type-(ii)" facts are those of the form "*I can* do *B* by means of doing *A*", i.e., the now familiar "doing *A* → doing *B*." "Type-(iii)" facts are worldly facts "concerning, e.g., the properties of figures or the location of bakeries" that "underwrite facts of types (i) and (ii)" (274; original emphasis throughout). (10) In classifying type-(iii) facts as he does, Small allows that they can take any number of forms "e.g. if *p* were the case, *q* would be the case; if an event of type *E* happens, an event of type *F* will happen; *F*s are *G*s; *Fa*, etc." (274). That is, type-(iii) facts are, for Small, analogous to the objects of speculative knowledge. Thus, they are anything that the agent can theoretically reason about. The domain of type-(iii) facts is arguably total.

(11) Small claims that facts of type-(i) depend on there being or having been at least one person for whom the fact in question is or was known as a type-(ii) fact (275). I take it that Small means to thereby distinguish facts about what one *could do* in principle or in such-and-such circumstances from type-(i) facts.

The former type of fact presupposes that the agent does not already have means sufficient to doing whatever it is she *could* do. These are type-(iii) facts. By contrast, type-(i) facts are about someone—possibly oneself—having a sufficient means to doing *B* in virtue of their ability to do *A*. We can think of the relationship between fact-types in terms of containment: type-(iii) facts contain both type-(i) and type-(ii) facts by virtue of underwriting them. And, in affirming the claim that what the agent knows in the exercise of practical knowledge is an event in the world as instantiating a particular means-end order, Small, following Anscombe, commits to the worldliness of type-(ii) facts. Every type-(i) fact is or at some point was, for some agent, contained in the corresponding type-(ii) fact but may not be contained in any type-(ii) fact for anyone else. And every type-(ii) fact brings with it a corresponding type-(i) fact for free. Thus, for most agents, a series of three concentric circles captures the relationship among the fact-types that they know:

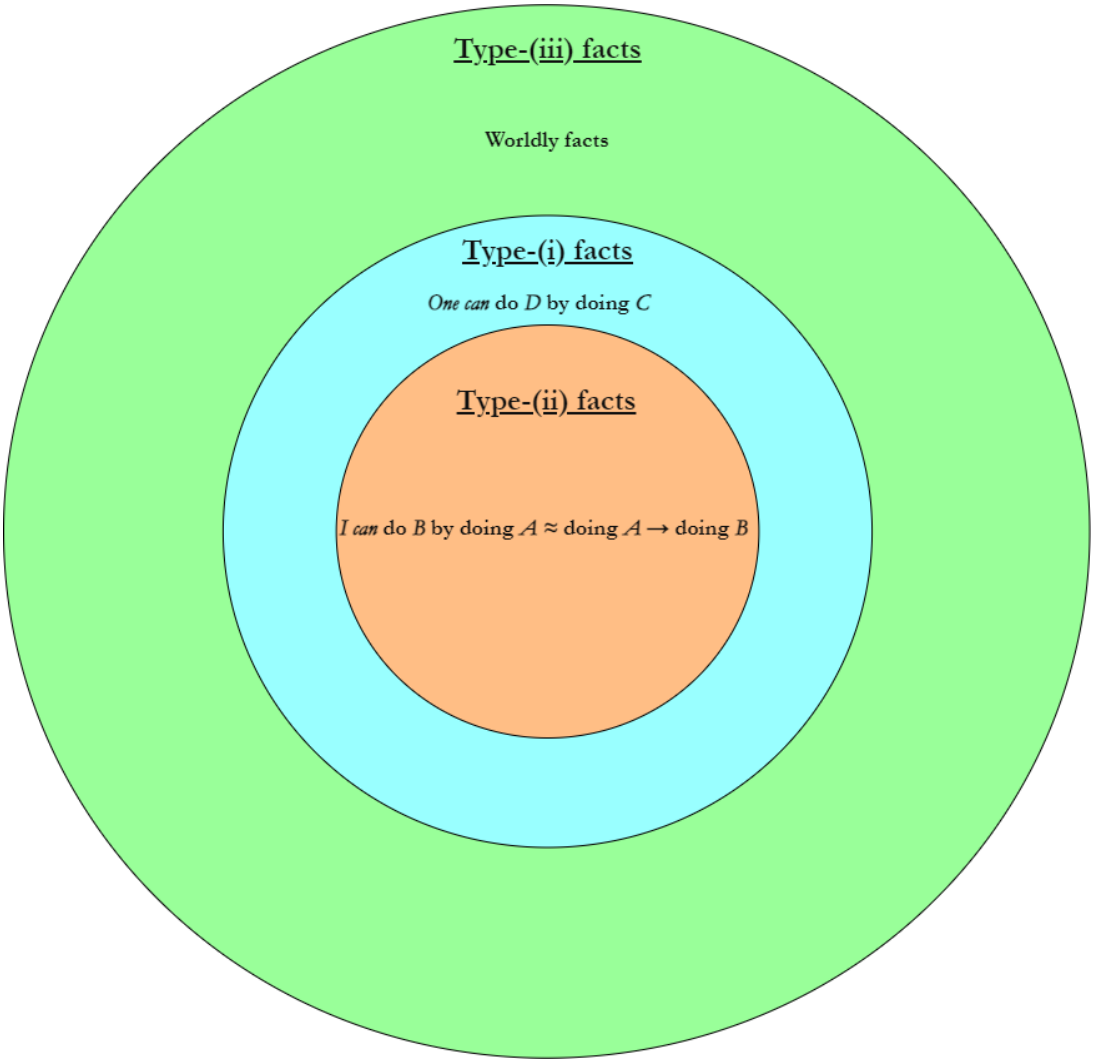

Figure B.1 Venn Diagram of Types of Facts as Known by a Particular Agent

Finally (12) Small begins to depart from Anscombe (2005) at this point. Anscombe claims that facts of the form "if plants are fed with certain substances, there will be spectacular plant growth," "if these substances are in the soil, the plants will be fed with them," etc. can feature in practical reasoning, in reasoning towards "a hypothesis proposed for investigation," and in the theoretical reasoning that is familiar to much of philosophy (126-127). That is, there is in Anscombe's thought the idea of, as Small puts it, "a logical kernel that can be used in [at least] two ways, but which is neutral between those two uses" (2022: 275). This kernel would consist of type-(iii) facts. By contrast, he claims that, given the form of practical reason(ing) as he has developed it, such facts "can be used in theoretical reasoning, but have no place in practical reasoning" (275). That is, if practical reason(ing) just is connecting series of intentional action concepts to each other with the means-connective, starting from some end which is not practicable to an action that is, then there is no room in such reasoning for facts other than "type-(ii) facts (and their type-(i) counterparts)" (275). Thus, there is no room in such reason(ing) for type-(iii) facts. Rather, type-(iii) facts support practical reason(ing) by providing agents with speculative knowledge of the "topic of an intention" or of "exactly what is happening at a given moment (say) to the material one is working on" (Anscombe 1957/2000: §28 & §48, respectively). Knowledge of type-(iii) facts is surely indispensable for bringing things off (see Appendices C-E). But they are not what an agent has practical knowledge of nor what she practically reasons about (*cf.* McDowell, 2022).

On one way of reading this last elaboration of Small, what he says must be false: if type-(ii) facts are also worldly by being made true by the agent then there are at least some type-(iii) facts about which the agent practically reasons. Insofar as she is now doing *B* (by doing *A*) or going to do *B* (by doing *A*) and insofar as she can calculate further using this fact, she is practically reasoning about type-(iii) facts. On another way of reading this last elaboration, what he says is arguably true: the agent's conception of the relevant type-(ii) fact(s) is fundamentally different from that of the corresponding type-(iii) fact(s). Were she to consider the relevant type-(ii) fact purely as a type-(iii) fact she would not be practically reasoning about it. And this is because such facts do not essentially involve her as agent. That is, type-(iii) facts, by virtue of having a fully general form such that their domain is anything knowable by her, are not essentially first-personal. But practical reason(ing) is essentially first-personal and, as such, takes as its material facts that are essentially perspectival, that is, that have an essential *de se* mode of presentation. The way out of the muddle is to recognize that a subset of worldly facts is characterized by its being possible for such facts to have such a mode of presentation and that only the agent, by bringing them about in acting, relates to those facts by conceiving of them as type-(ii) facts. Yet, she can likewise relate to them as type-(iii) facts. When she does both, "we can really speak of two knowledges" of one

object (Anscombe 1957/2000: §48). Her relating to them as type-(ii) facts can be expressed by an "account that [she] could give of what [she] was doing, without adverting to observation" (§48). Her relating to them as type-(iii) facts can be expressed by an "account of exactly what is happening at a given moment (say) to the material [she] is working on" (§48). There is therefore a logical kernel that is shared between theoretical and practical reason. And this kernel consists in the intentional action concepts that the agent possesses, draws on in instantiating the corresponding intentional actions and applies in deliberation such that she relates to them under the *de se* mode of presentation, or applies in perceptual judgment or theoretical reasoning such that she relates to them as type-(iii) facts. Anscombe is right about this much. What Small is right to point out is that the kernel is not neutral between theoretical and practical reason. Rather, as I argued in chapters 1 and 2, the kernel has its ground in practical reason and becomes an object for theoretical reason only by virtue of some agent's making it true.

So, altogether, practical reason is in the first instance a domain-specific capacity for ordering intentional action concepts that the agent possesses. The form of its activity is stringing together such concepts through the means-connective from some future contingent affectable by action to some practicable means (relative to the agent's skills) that the agent (perceptually) recognizes as a "*this* that I can do." The intentional action concepts strung together are, for the agent, type-(ii) facts of means-end orders which she (practically) knows she can bring about. Type-(i) facts have a role in the rational transmission of skill and type-(iii) facts aid practical reason by providing material for (perceptual) recognition with a view to further calculation and further action. Yet, the core of practical reason is the capacity to calculate means to ends, the form of which has been more fully described in this Appendix. This capacity is at work both in deliberation and in knowledgeable intentional action in the absence of such deliberation. The former is a mental action of at least potentiating (/generating intentional action acorns for) the construction and use of discursive representations of ends to be achieved and means of achieving that end appropriately connected to each other by "by means of," that is, by the means-connective. The latter is the potentiation of already internally calculatively structured networks of intentional action concepts constitutive of (parts of) the agent's skills and habits in light both of the agent's practically wanting something and her (perceptually) recognizing the appropriate practicable means relative to the appropriate target in the context of action. Both deliberation and non-deliberative cogent practical cognizing exhibit the calculative structure that practical reasoning in general has the function to provide and which is constitutive of corresponding intentional actions.

Appendix C: *Prima Facie* Failures of Practical Cognition and Mistakes in Performance

Drawing from chapters 1 and 2, other things being equal and assuming no change of mind occurs, in cases of non-mental action, practical want either drives search for practicable means through mental acts of deliberation or through (perceptual) recognition and subsequent implementation of practicable means culminating in the performance of an intentional action. This is just the exercise of the agent's cogent practical cognition constitutive of an exercise of her practical knowledge. With the practicality of practical knowledge already on the table, what remains is to explain its intelligence and, in particular, how it achieves the status of *knowledge*. What makes practical knowledge *knowledge* and not mere belief, opinion, wanting, or even brute desire? The answer can be found by considering the factualist understanding of such knowledge more closely. In chapters 1 and 2, I discussed how speculative knowledge of type-(iii) facts contribute to the cogency of practical cognition as well as how successful performance non-accidentally makes it true that an agency is intentionally doing *B* (by means of doing *A*/in order to do *C*) *purely* in virtue of the success of her practical cognition. I will have more to say about non-accidental success in Appendices D and E. For now, I want to unpack and discuss types of *prima facie* failure specific to practical cognition as well as mistakes in performance. The relevant failures are *prima facie* because they can be overridden by the agent's correcting her behavior. Discussing them helps put the shape of cogent practical cognition culminating in the exercise of practical knowledge into relief. Along the way, I appeal to Small's (2012) discussion of the "cognitive condition" which amounts to the factualist understanding.

In any interaction between finite beings and their environments, the finite being's side of the interaction is limited and can fail to appropriately hook up to or match the environment. Practical cognition is no different: it too is fallible. Given this, it sounds odd to say that agents act intentionally only when they have pure practical knowledge of what they are doing, why they are doing it, and how they are doing it. Can not someone mean to do something and succeed despite being skeptical of their success? And cannot someone mean to do something and succeed despite being unjustified in thinking that they are doing it? To hold onto the factualist understanding of pure practical knowledge, I need to assuage these worries. And to start to assuage them, I need to consider the ways in which practical cognition can fail (at first) to achieve the status of pure practical knowledge and cases in which the agent's performance fails to live up to her otherwise cogent practical cognizing. Consideration of these limitations on success show the limits of practical cognition. In Appendices D and E, I show that apparent counterexamples to the factualist understanding fall along the edge of these limits. Only when the apparent counterexamples are situated relative to these limits will it be clear what the

factualist understanding really amounts to. It is not a flatfooted position. And only then will it become clear that these worries are misguided. So, let us start by staking out the limits themselves.

Practical cognition can *prima facie* fail in at least the following five ways:[130]

A) ***Futile Effort*:** something is not in one's power to do or bring about but one thinks that it is.

B) ***Practical Error:*** one thinks that one knows how to do something but one is mistaken.

C) ***Practical Undercutting:*** something that must be the case for one to succeed in doing *D* by means of doing *A* is not the case and one fails to have speculative knowledge of this fact.

D) ***Thwarted Attempt:*** one's intention is contradicted by another's.

E) ***St. Peter*:** one has (infallible) foresight that one will do *D* intentionally and yet intends not to do *D* but is ignorant as to the way in which they will end up doing it.

There are, as well, **mistakes in performance**, where one does something other than what they intend despite having fixed both on an end that they can do or bring about and on actually practicable means to that end. In cases of **futile effort**, **practical error**, and **practical undercutting**, the failure in question results in what Small (2012), taking some liberties from Anscombe, calls the agent's intention "falling to the ground" (146*ff*.; *cf*. Anscombe 1957/2000: §32). An intention falls to the ground whenever the agent is at the moment not satisfying her intention by whatever means she takes herself to be implementing because she fails to possess that means at that moment. By contrast, in cases of **mistake in performance**, the agent is in possession of the relevant means and is simply failing to implement it. Her practical cognition is perfectly cogent: were she taking the means she has settled on, she would be, e.g., doing *D* by doing *A*. Yet, what she would say she is doing is directly falsified by her performance (Anscombe 1957/2000: §32). Such a thing can happen when the development of a potentiated solution system into an ongoing performance is defective due to vagaries in the environment or in the agent. (Such vagaries stunt the growth of the intentional action "oak.") In cases of **thwarted attempt**, the agent's intention is contradicted by another intention. And, in **St. Peter**, the agent has pure practical knowledge of what she is doing under one description, e.g., doing *C*, which fits with what she is seeking to avoid, e.g., doing *D*, but is not mindful of or is self-deceived that her intentional action of doing *C*

[130] I take some of the naming conventions from Schwenkler (2019: 181-182).

just is an intentional action of doing *D* (Schwenkler 2019: 207-210). This is a kind of defective practical knowledge. I will not discuss **St. Peter** further, since the other *prima facie* failures and **mistakes in performance** are together more instructive for grasping how practical knowledge is knowledge by instantiating an objective order and, thereby, making its object true. The kind of failure exhibited in **St. Peter** is more instructive for understanding certain aspects of moral self-knowledge, of which pure practical knowledge might well be a species.

Starting with the *prima facie* failures of practical cognition that result in an intention's falling to the ground, let us consider **futile effort**. Consider an agent who practically wants to attend class on what turns out to be a holiday that her university observes. The agent thinks that a sufficient means to attending class is going to the location where class is normally held at the appropriate time. She is not wrong because this is not a way of attending class on any normal class day (unlike **practical errors**), because there is some type-(iii) fact that makes successful performance in this case depend on taking what is for her some unusual or exceptional means (unlike some cases of **practical undercutting**), because she simply does other than she intends (unlike **mistakes in performance**), because she is prevented by someone intent on stopping her (unlike **thwarted attempts**), or because she is conflicted or self-deceived (unlike **St. Peter**). Rather, the agent is just wrong about what she can achieve by going to the relevant location at the relevant time. Her going there is a **futile effort** to attend class. On that day, she cannot satisfy her intention to go to class. However, if her intention is just to go to class then she can satisfy *that very intention* come the following week.

Moving on to **practical error**. Consider an agent who practically wants to tie a full Windsor knot and who takes himself to know how to do this but does not (yet) know to do a second loop with the wide end of the tie on his left side. Anytime the agent tries to tie a full Windsor, he ends up with a half Windsor. He is not wrong because at present tying a full Windsor is not within his power to do or bring about (unlike **futile effort**) or so on.[131] Rather, the agent is just wrong about what it takes to tie a full Windsor. He has made a **practical error** in attempting to tie the full Windsor. Importantly, it is the same **practical error** he would have made had he been deliberating about how to tie a full Windsor and judged that he could do so by going through the following steps:

[131] I leave showing how **practical error** is not reducible to **mistake in performance**, **thwarted attempt**, or **St. Peter** and how it does not exhaust the category of **practical undercutting** as an exercise for the reader.

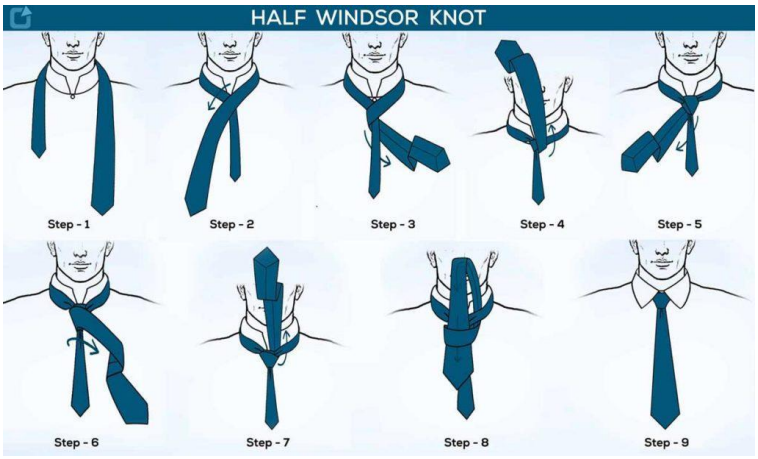

Figure C.1 Half-Windsor Knot

A pictorial representation of the steps involved in tying a half Windsor knot. Taken from nexoeye.com: https://nexoye.com/half-windsor-knot/

Taking what he takes to be the means to tying a full Windsor will leave his intention to tie one unsatisfied. However, he can satisfy *that very intention* once he is made aware of his mistake, goes on to search for or calculate the actual practicable means, and implements that means.[132] What the agent lacks in this case and what is required to have the skill to tie a full Windsor is making an ability to form a second loop determinate part of the shape *in concreto* his skill for tying ties. His acquiring the skill to tie a full Windsor might come with his acquiring this latter skill for tying ties, integrating the relevant parts of his skill, say, to tie knots (or his shoes) with these newly acquired skills. Or his acquiring the skill to tie full Windsors might simply be a matter of adding to the determinate shape *in concreto* of an already existing skill of his to tie ties, integrating parts of the latter with this newly acquired skill. He might likewise incorporate the skill to make the second loop or his new skill for tying full Windsors into corresponding habits, making either into a constitutive unity that he can instantiate as a basic action.

Moving on to **practical undercutting**. Thus far, I have considered *prima facie* failures of practical cognition that have their grounds in interesting subsets of a more general kind of *prima facie* failure, namely, **practical undercutting**. Indeed, in cases of **futile effort** and **practical error**, the agent's failure is due to ignorance about what would have to be the case for them to succeed in doing what they are now doing. They lack relevant bits of speculative knowledge of type-(iii) and type-(i) facts, respectively, which would undercut

[132] We will see later, in Appendices D and E, how this differs from cases in which the agent has to learn how to do something in order to be able to implement the relevant means. In the latter cases, the agent merely *aspires* to do whatever action requires his learning the relevant means (*cf*. Small 2012: 171-172).

the intention, e.g., to go to class or that one can manipulate the tie thus-and-so in order to tie a full Windsor. And every instance of **practical undercutting** is such that the agent's current efforts are made futile and her judgment as to a practicable means are incorrect on the relevant occasion. That said, there are cases of **practical undercutting** that stem from unforeseen circumstances and bad luck. The case that Anscombe (1957/2000: §23, 31-32) considers in introducing the idea of what Small call's an "intention's falling to the ground" is a case in point: a man attempting to pump water into a cistern is not doing so due to there being a hole in the pipe running from the water source to the cistern. Assuming that it is within his power to plug the hole or find an alternative means, his goal of replenishing the water supply is still within reach for as long as he maintains his intention. And assuming the hole is not a design feature of the pipe, his judgment regarding what constitutes a sufficient practicable means of achieving his goal is incorrect only because of the circumstances or because such circumstances are unlucky. He is not confused about how to do his job. He is simply ignorant of a development in the circumstances of action that complicates that job. If he keeps trying to use the pump as is, his intention will go unsatisfied. However, he can satisfy *that very intention* once he is made aware of the new circumstances by plugging the hole or, say, going to fetch some pales (Small 2012: 157-158).

Before moving on to **thwarted attempt**, I want to consider **mistakes in performance** and the so-called "broadness" and "openness" of progressive judgments. Focusing on **mistakes in performance**, such mistakes are not instances in which an agent's intention "falls to the ground," because in such cases the agent is in possession of the appropriate means. She is simply failing to implement those means by doing what she is doing. Consider an agent who practically wants to reach the top floor of a building and knows that she can get to the top floor by riding the elevator and can ride the elevator by calling it and can call it by pushing button *M*. And suppose that pushing *M* is for her a "*this* that I can do." Finally, suppose that, nonetheless, she ends up pressing button *N* by mistake. This agent has an appropriate object of practical want. She has correct knowledge how to call and ride elevators. She is not unaware of any facts that are in this case pertinent to her getting to the top floor by means of taking the elevator, taking the elevator by means of calling it, or calling it by means of pressing button *M*. She is not being stopped by someone, say, forcing her hand towards button *N*. And she is not conflicted or self-deceived about what she is doing. Her performance simply fails to match her cogent practical cognition, that is, her intention. In this case, her performance is a "*direct* falsification" of her intention to press button *M* (Anscombe 1957/2000: §32). The falsification is direct because there is no mediation from incorrect judgment on her part, from unwieldy circumstances or bad luck, nor from someone else's intentionally impeding her to the failure to do what she intends (*cf.* McDowell 2022). That is, to quote

Anscombe quoting the author of *Magna Moralia* "the mistake is not one of judgment but of performance" (1957/2000: §32). Indeed, it is important that in setting the stage for **mistakes in performance**, Anscombe emphasizes that the relevant situation is one in which "this and this alone [i.e. the performance itself] constitutes a mistake" (§32). For only in such cases is it true that the agent's practical cognizing is nonetheless cogent. Our finitude is such that such mistakes are a possibility: even when the agent's practical cognizing is in perfect working order, her performance can sometimes be other than it should. In the case under consideration, if the agent does not go on to do anything else, her intention to go to the top floor will go unsatisfied. However, she can satisfy *that very intention* by correcting her performance, that is, going on to actually press button *M*.

So far, I have just considered *prima facie* failures of practical cognition that result in an agent's intention "falling to the ground" and **mistakes in performance**. Again, in the first set of cases, such failures just mean that the agent is not at the moment satisfying her intention by whatever means she takes herself to be implementing because she fails to be in possession of the appropriate means. In **mistakes in performance**, the failure is due to defective development from potentiation to ongoing performance. Note that in each case of *prima facie* failure, the agent can "pick [that very intention] back up again" (Small 2012: 147). That is, in each case, the agent can go on to satisfy the very intention that she is at the moment failing to satisfy by coming into possession and implementing the appropriate means. In cases of **mistakes in performance**, she need only become aware of the mistake and implement the appropriate means of which she is already in possession. As Small points out, an agent's ability to pick her intention back up and correct for mistakes is connected to the so-called "broadness" and "openness" of progressive judgments. Let's take a short detour through these notions.

Starting with broadness, following Thompson (2008), Small points out that judgments of the form "'*S* is[(/was)] doing *B*' might be true even if there is nothing that *S* is doing right now that amounts to doing *B* or taking some means to doing *B*" (2012: 177). Such judgments are broad in that their truth depends on *S* being engaged in an ongoing activity or in a social praxis. With respect to activities, it is true of, say, a duck that it is flying south for the winter even if, at the moment, it is dipping its head below the surface of a pond looking for a bite to eat. With respect to social praxes, it is true of a person that, say, she is baking a loaf of bread even if, at the moment, she is lying on the couch. After all, in such a case, she might be waiting on the dough to rise or on the bread to bake. Because the powers framework understands actions as teleologically unified developing processes, for a happening to count as (part of) an action does not require that the agent is at all times during that happening causing the action to come about. For it can be true of the agent in cases of **futile effort**, **practical error**, **practical undercutting**, and **mistake in performance** that she is doing what she intends

even as she is failing to do it. Returning to our examples, whether it is true that the agent in question is performing the relevant action depends on the student's respecifying her intention as to have the relevant object which is in her power to do or bring about, the "tie-tyer" correcting his misconception about how to do the relevant thing, the pumper addressing the development in the circumstances of his action, and the "elevator-caller" correcting her performance, respectively. Other things being equal and so long as the agent does not change her mind and goes on to do whatever she needs to so as to pick her intention back up or correct her performance, her practical cognition ultimately succeeds in its aim and thereby achieves the status of an exercise of practical knowledge. In which case, it was true *all along* that she was doing *B* and doing so intentionally.

Before moving on to the openness of progressive judgments, there are two wrinkles to consider. First, in at least some cases of *prima facie* failure or **mistake in performance**, the agent might need to respecify her intention. Second, in at least some such cases, the only description of the action for which it is true *all along* that the agent was performing it might well be the one describing her achieving her end. On respecifying one's intention, consider a variation of Anscombe's pumper case that Small presents: "suppose that [the pumper's] stroke is too shallow; it does not draw the water, and in fact he is only pushing air through the pipe" (2012: 157). In such a case, the pumper is making a **mistake in performance**: "there is nothing wrong with his *intention* (to replenish the water supply by pumping); it is his *execution* (his pumping) that is at fault" (157; original emphasis). Suppose further that he is asked how he is pumping, and he expresses a further articulation of his intention: "to replenish the water supply by pumping, by making *these* movements with [my] arm" (157; original emphasis). In such a case, in order to correct his performance, the pumper has to respecify his intention, since "making *those* movements with his arm is not a way of pumping" (157; original emphasis). That is, in adopting a very slightly different course of action, while it is true all along that he is replenishing the water supply and doing so by pumping, it is only because he goes on to make movements that are different from "*these* movements" that this is so. This is a function of his practical cognition: he holds his end fixed and searches for an alternative practicable means.

As Small points out, this first wrinkle is instructive:

> Examples such as these illustrate what is invariably the case: that there is a kind of constant correction that goes on throughout the course of an intentional action, as the agent responds to the miniature successes and failures, obstacles and alternative possibilities both foreseen and unforeseen, that he encounters in what he is doing and what he is acting on and with. What I am bringing out is that this constant correction amounts to *rationally further determining or respecifying one's intention*: the answer to the question 'How?' is constantly being finessed (158; original emphasis).

This leads to the second wrinkle: so much correction or finessing of an intentional action might go on that only the description of the agent's achieving her goal counts as having been true *all along*. Another variation on the pumper example brings this out (146-147). Suppose that there is a hole in the pipe and that the pumper cannot fix it. Small gives what reads as a plausible response to the situation from the pumper's employer: "The pipe cannot be repaired, so the intention to *replenish the water-supply by pumping water through the pipe* cannot be executed [...] but the inhabitants need running water and there are two buckets over there—get to it!" (147; original emphasis). Assuming the pumper gets to it, the only description of what he was doing and of which he had practical knowledge of *all along* was replenishing the water supply. In which case, it was true *all along* that he was doing so and doing so intentionally.[133]

The point of the two wrinkles is that the transition from *prima facie* failure of practical cognition or **mistake in performance** to outright failure is not a matter of course but, rather, requires more than that the agent is at present failing to satisfy her intention. We can see this even more clearly by considering the openness of progressive judgments. Openness has to do with the entailments of judgments of the form "*S* is/was doing *B*" for some specific *B*s. As Small puts it: "while the truth (at t_1) of '*S is doing B*' entails (at t_2) the truth of '*S was doing B*', [for some *Bs*] neither entail—at any time—the truth of '*S did B*' [for those *B*s]" (177; original emphasis. See also Anscombe 1957/2000: §23). A tired example: one can be walking across a street, can have been walking across a street, and, nonetheless, never walked across the street, say, because he gets hit by a bus. Small argues that the relevant subset of intentional action concepts subject to openness admit "a contrast between two ways in which something that is doing *B* can cease to be doing so: by *stopping* doing *B* or by *finishing* doing *B*" (2012: 177; original emphasis). The event predicates that fill the role of *B* and which admit of this contrast are, following Vendler (1957), Kenny (1963), Comrie (1976), Galton (1984), Rothstein (2004), and Small (2012), *telic*. Event predicates that do not admit of the contrast between stopping and finishing are *atelic*.[134]

---

[133] That said, the pumper will also have practical knowledge of whatever means he ultimately successfully implements, be it plugging the hole, lugging pales, or whatever. The point here is just that *what* he can be said to have practical knowledge of *throughout the whole process of correction and respecification* might be no more than "replenishing the water supply."

[134] Atelic event predicates do not specify an end. Walking and growing are atelic event types. Note that one can transform an atelic predicate into a telic one by specifying an end, e.g., walking to work or growing up. As Small puts it "by specifying an end, one specifies the proper terminus of the event-form to which the predicate refers itself, whose instantiation on an occasion may then be subject to proper termination" (2012: 176). On the powers framework, all intentional actions are telic. However, some actions taken as means may well be atelic, for instance, in cases of walking to a place or taking a walk. In the latter case, "going for a walk" is telic, even if what constitutes the "end" in such a case is highly ambiguous, while "walking" in such a case is atelic.

Telic event predicates specify an end such that their instances can be stopped by being interrupted.[135] This suggests that telic events have a "proper terminus" or endpoint at which they aim (Small 2012: 178). Of the telic events that are intentional actions, then, the interruption of those actions and subsequent failure to pick the intention back up or correct performance results in an incomplete action. But note that this means that, even in such cases, although the agent did not do *B*, it is nonetheless true that the agent *was* doing *B* intentionally at least up until the interruption. And this means that, at least up until the interruption, the agent's practical cognition achieves the status of an exercise of practical knowledge of her doing *B*. Moreover, for those actions that are *atelic*, that the agent *is* or *was doing B* entails that she *did B* for as long as she is or was doing *B*. Atelic intentional actions like walking, running, pumping, typing, etc. are such that even if the intention with which the agent does these things falls down or she makes a **mistake in performance**, so long as she at any point was doing them intentionally then she counts as having done them intentionally at those times. Therefore, at those times, she counts as having and exercising practical knowledge of both doing and having done them. Thus, at those times, she does them and does them intentionally,

The ability to pick an intention back up and correct for **mistakes in performance** as well as the broadness and openness of progressive judgments altogether suggest that outright failures of practical cognition are the rare exception. Other things being equal and assuming no change of mind occurs, in order for her practical cognition to outright fail to achieve the status of an exercise of practical knowledge, the agent must fail to act as she intends in cases where there is no (more) opportunity for correction or where she does not check whether she succeeds when success is not a matter of course. In either case, she does not exercise practical knowledge of what she is aiming to do and, thus, if she does them at all, is not doing those things intentionally. Nevertheless, it is worth emphasizing that even in these cases, when they concern actions that are non-basic at least, the agent will still exercise practical knowledge of her action "in its more immediate descriptions" (Anscombe 1957/2000: §48).

The only possible exception concerns *prima facie* failures of practical cognition or **mistakes in performance** attaching to an agent's immediately practicable means themselves. For one possible case, see Setiya (2008: 389-390). In this case, Setiya's hand is numbed and was paralyzed. In a moment of practical irrationality, he intends to clench his fist and happens to do so because the paralysis has worn off. Setiya assumes that his numbed fist clenching is intentional. Yet, he acknowledges that his being in a state of intending to

[135] Such events include non-actions: trees falling or dropping their fruits are telic and can be interrupted but trees do not perform actions.

clench his fist and believing he can do so is irrational. In which case, I find his assumption that his fist clenching is intentional highly dubious. To my ear, the case strikes me as analogous to someone who has never played darts intending to make a bullseye by making a trick shot and believing he can do this (see also Appendix D). If he succeeds, it is purely by luck and not by virtue of his fixing on and implementing what are practicable means for him. I would thus not classify his making the bullseye by making the trick shot as an intentional action of his: it does not instantiate a means-end order that he calculates. Rather, he *aspires* to hit the bullseye by making the trick shot and intentionally makes movements which he takes to approximate the trick shot. My sense is that Setiya's case does not itself decide between our warring intuitions. I will have more to say about aspiration in Appendices D and E.

Before concluding this Appendix, it is worth pausing for a moment to consider cases of **thwarted attempt**. Suppose the pumper intends to replenish the water supply and someone else expresses a contrary intention, "Oh no you will not!" thereupon setting out to make a hole in the pipe with a pickaxe. In this case, assuming the second agent fixes upon some practicable means, unless some measures are taken by someone to guard against her satisfying her intention, the pumper will fail to have or exercise practical knowledge of replenishing the water supply. Notice that, unlike when an intention falls to the ground or when the agent makes a **mistake in performance**, the pumper can no longer simply hold his end fixed and correct his behavior or a misconception either about what is within his power or about how to do something. Rather, assuming that the "pipe-perforator" is intent on preventing him from replenishing the water supply, the pumper has to find some practical means of stopping her. This is because the logical relation between the two intentions is different than that between the relevant undercutting fact or **mistake in performance** and the pumper's intention. Namely, the pipe-perforator contradicts the pumper. If he does not form an intention to stop her from stopping him and both parties remain resolutely at odds, his intention to replenish the water supply will go unsatisfied. It is only by someone's successfully contradicting her with a yet further intention, say, to station guards along the pipe to prevent her from perforating it, that the pumper can retain and exercise practical knowledge of replenishing the water supply and, so, is replenishing the water supply and doing so intentionally.[136]

[136] Note the difference between this and the following case: suppose the pipe-perforator succeeds at creating a hole in the pipe and thinks this alone is sufficient to stop the pumper or thereafter changes her mind about preventing him from replenishing the water supply. If the pumper is unaware of this development then he is back in the **practical undercutting** situation. This implies that contradiction among intentions lasts only so long as each contradictory is maintained by the respective agent.

In this Appendix, I considered the ways in which practical cognition threatens to fall short of practical knowledge. The three general categories I have outlined are **practical undercutting**, **mistakes in performance**, and **thwarted attempt**. And these considerations have led me to the following three conditions for outright failure of practical cognition:

> To outright fail to have or exercise practical knowledge of her doing *A*, the agent must intend to do *A* and must fail to do *A* where
>
> > i) there is no (more) opportunity for correction;
> >
> > ii) she does not check whether she succeeds in doing *A* where success in doing *A* is not a matter of course; or
> >
> > iii) no one contradicts an intention that contradicts her own.

These conditions suggest that cogent practical cognizing is robust and flexible in the sense that it tolerates *prima facie* failures, **mistakes in performance**, respecifying, and correction or finessing. Moreover, they suggest a shape for a positive account of how practical cognition achieves the status of an exercise of practical knowledge on the factualist understanding of such knowledge. Practical cognizing must allow for respecification and correction throughout the course of the performance of some action. Moreover, practical cognizing must draw on speculative knowledge (or, at least belief or opinion) of "the topic of an intention" or of "exactly what is happening at a given moment (say) to the material one is working on," especially in cases where success is not a matter of course (Anscombe 1957/2000: §28 & §48, respectively). As I argued in chapters 1 and 2, the dependence of practical knowledge on speculative knowledge is not of a kind on which one has practical knowledge of doing *A* in virtue of having speculative knowledge of doing *A*. Rather, as I argued, her practical knowledge is already present in any case where she gains speculative knowledge of her doing *A*. The role such speculative knowledge plays is as an aid in the production of further practical knowledge.

Appendix D: Practical Knowledge, the Carbon Copier, and Pujols

In this Appendix, I present a more thoroughgoing positive account of practical knowledge adopted by the powers framework that fits its factualist understanding. I have already identified the capacity which lies at the heart of this discussion, namely, the capacity for means-end calculation or—what amounts to the same—practical reason (Sections 1.3.2-1.3.4). Practical cognition consists in the activity of this capacity for calculation and its *generans*, namely, practical want. The capacity for means-end calculation is itself a complex rational capacity, one that becomes a determinate part of the shape *in concreto* of *every* skill and is incorporated into every *intelligent* habit.[137] Intentions, then, are nothing other than the activation of the relevant part of the respective skill or habit. That is, intentions are just potentiations of the intentional actions into which they develop. They are intentional action acorns. So, on the powers framework, intentions are not mental events that exist independently of their objects and stand to them as isolable truth-conditional representations which somehow (non-deviantly) bring those objects about. Intentions are not a required "extra feature" that structure action by (merely efficiently) causing it (*cf.* Anscombe 1957/2000: §48). Put another way, they are the exercises of the capacity for practical reason that order the exercises of skills and intelligent habits into means (and ends) and *are nothing other than such ordering*. They are the activation of skills or intelligent habits that represent their objects—the exercise of such skills or habits—by virtue of being a constitutive moment of the teleological unity of the development of those objects. They represent by standing for the actions of which they are potentiations as appropriate for instantiation in the relevant context given the presence of the relevant target(s) and need(s). The capacity for means-end calculation itself has a shape consisting in discursive skills for inferentially and perceptually identifying actions (and goals) as such within some domain and for connecting them to each other as means (and ends). Moreover, the capacity can subsume habits the agent has developed for such recognition and stringing together, that is, whatever calculative heuristics she uses in pulling off actions in the relevant domain.

The intelligence of the calculative capacity consists in the agent's flexibility and non-accidental success regarding i) taking as objects future contingents affectable by her action, ii) ordering intentional action concepts from those future contingents to some appropriate practicable means, and iii) implementing those practicable means in intentional action. The successful exercise of this capacity is the exercise of practical knowledge, i.e.,

[137] Importantly, on the powers framework, practical reason is not a capacity that can exist in isolation of skills or habits for other actions. One cannot have practical reason in the absence of possessing any other skills or habits.

the making true of an agent's doing *D* by means of doing *A* vis-a-vis *this*. This leaves out how practical reason avoids settling on impracticable means, how practical want avoids settling on unattainable ends, and how performance avoids failing to live up to the agent's otherwise cogent practical cognizing. I discussed this topic in the previous Appendix as well as Section 1.3.4. There I was focused on the robust flexibility of practical cognizing, how the calculative structure it gives to intentional actions underwrite those actions' temporal structure, and how speculative knowledge aids practical knowledge. But beyond cogent practical cognizing considered by itself and with the aid of speculative knowledge, what allows the agent to latch onto an appropriate end and onto an appropriate means such that she knows "that it will be no accident if [she] ends up having intentionally done what [she] is doing intentionally (or intends to do)" (Small 2012: 136). Following Small (2012), the rest of the answer lies in distinguishing between cases where an agent intends to do something she is not yet skilled at and cases where she merely "aspires" to do such a thing (167*ff*.) This part of the answer helps address apparent counterexamples like the carbon-copier.

Let's drill down on such counterexamples. Cases like the carbon-copier are those in which an agent means to do something and supposedly does it unknowingly. Failures of the sort discussed in the previous Appendix are in a way uninteresting, since everyone agrees that the relevant action is not being performed anyhow. What is interesting is unwitting success. To address cases of unwitting success, I take the distinction between intention and aspiration from Small (2012) who himself takes it from Baier (1970). This is how Small presents the idea, focusing on the fact that what makes the difference is procedural knowledge-how of *some* means or the lack thereof:

> One can intend to do *A* without knowing how one *shall* do *A* so long as one knows how one *can* do *A*. But if one does not know how one *can* do *A,* then one can merely aspire to do *A* (where the inner tendency of this aspiration is towards the acquisition of knowledge how to do *A*, and thus towards its own transformation from aspiration into intention) (2012: 170; original emphasis).

Thus, one who aspires to do *A* cannot actively synthesize or potentiate doing *A*, where her ability to do either of these things is the mark of (because constitutive of) her being able to intend to do it. Taking an example from Setiya (2008) and Paul (2009), an agent can only aspire to dance the tango at her wedding so long as she does not know how *she* can tango. By contrast, an agent who knows how she can dance the tango can intend to do it even if she does not know how she shall do it, say, because she has not yet identified a suitable dancing partner for the relevant occasion.

This distinction suggests distinct gaps in practical cognizing, specifically with respect to its role in the agent's calculating means to her ends. Whereas intending admits of gaps in the calculative order that the agent synthesizes due to a lack of the relevant type-(iii) facts which would underwrite her action, it does not admit of gaps due to a lack of the relevant type-(ii) facts. The kind of gaps that practical cognition tolerates versus those that it does not again brings to the fore the role of speculative knowledge in supporting cogent practical cognizing. Agents make use of evidence in coming to possess speculative knowledge of type-(iii) facts which help them fill in the details of their calculations. Our dancer identifies a dancing partner through perception, inference, testimony, etc. But this does not mean that her practical knowledge of dancing the tango with that partner is based on perception, inference, testimony, etc. Rather, her practical knowledge thereof stems from her cogent practical cognizing. In this case, she recognizes dancing thus-and-so with this person vis-à-vis *this* as a practicable means to dancing the tango through perception. The activity of cogent practical cognition enters the scene and issues in action when she actively synthesizes or immediately instantiates (through skillful action) the order of dancing the tango by dancing thus-and-so with this person vis-à-vis *this*.

Setiya and Paul might be novice dancers who do not know how to dance the tango. This means that they lack procedural knowledge-how of the relevant type-(ii) facts that our dancer has. As such, they can merely aspire to dance the tango. Aspiration is nonetheless a conative attitude: it engages practical cognition by directing the agent to calculate means to coming to possess the relevant procedural knowledge. On the powers framework, this amounts to saying that recognition that one fails to possess the intentional action concepts needed to perform the relevant action potentiates exploratory actions for coming to possess those concepts, or what amounts to the same thing, those means. That is, such "practical conflict" as Shepherd (2023) calls it between wanting to do something and recognizing that one fails to have the means to do it tends to initiate deliberation towards acquiring those means. In the case under consideration, other things being equal and if no change of mind occurs, if they aspire to dance the tango then Setiya and Paul will deliberate what to do. In all likelihood, they will settle on going to dancing classes, practicing movements that they learn in class, and attempting increasingly sophisticated sequences, forming corresponding intentions. Alternatively, they might settle on attending classes only long enough to be able to learn to actively synthesize and then instantiate movements constitutive of one dance to be performed, say, at their wedding reception. In this latter case, their aspiration is better cast as dancing the tango once at their wedding. What the foregoing example illustrates is, as Small notes, that aspirations have an inner tendency towards becoming intentions themselves once the agent has acquired procedural knowledge-how of the relevant type-(ii) facts. That is, aspirations trend towards becoming

intentions once the agent has acquired the relevant intentional action concepts, when she comes to possess the relevant means. And aspirations do this by the agent's engaging in cogent practical cognizing the means to acquiring that procedural knowledge either in the form of a synthesis of intentional action concepts in calculation (which concepts ultimately bottom out in immediately instantiable habits) or in the form of manifesting some appropriate acquired skill.[138]

How does the distinction between intention and aspiration help us deal with the carbon-copier? To start, let us consider the original example from Davidson (1980/2001: 50, 82). Suppose an agent takes 11 sheets of paper to be marked and 10 sheets of carbon paper, interlaces the two stacks, and attempts to mark the top sheet with enough force that the carbon paper marks all of the other sheets. The idea is that the agent can be intentionally making 10 carbon copies in marking the top sheet without knowing that she is doing this. If this were right, then the factualist understanding of practical knowledge would be false. Now, Davidson leaves the case under described. So, first, recalling the previous Appendix, there is no reason to think that there is no (more) opportunity for correction, that is, that the agent will be kept from making a second attempt if she fails or that her intention is specifically to make 10 carbon copies all at once on the first try. If she does fail and then goes on to finish making the rest of the copies with less (and less) interleaved sheets, then it will have been true *all along* that she had practical knowledge of making 10 carbon copies and, so, that she was doing so and doing so intentionally.

Second, suppose that the agent checks after her first attempt to see if she has marked the 10^{th} copy. After all, making 10 carbon copies this way is difficult and success here is not a matter of course. There are two possibilities. Either she sees that she has succeeded or she sees that she has failed. In the first case, what the agent gets is speculative knowledge of having marked the 10^{th} copy. Her practical knowledge is of making 10 carbon copies. What is left at issue here is whether her practical knowledge of making 10 carbon copies depends on or is constituted by this speculative knowledge and, if so, how. If her practical knowing is to be constitutive of her doing then it cannot be that her practical knowing is justified or constituted by her perceiving that she has done it. Rather, following Section 1.3.4, her perceiving that she has marked the 10^{th} copy potentiates further actions. In the second case, what the agent gets is speculative knowledge of having failed to mark further than the n^{th} copy (where $n<10$). In which case, other things being equal and assuming she does not change her mind, she sets

[138] Thus, if the tie-tyer does not know how to make the second loop then he can only aspire to tie a full Windsor. If he does know how to make the second loop but does not know to do so then, if he is made aware of his **practical error**, he can intend to tie a full Windsor by bringing the relevant intentional action concept into his calculation.

about finishing the remaining copies. That is, we have returned to and elaborated the first variation of the case just considered. Again, assuming she goes on to finish them then it will have been true *all along* that she had practical knowledge of making 10 carbon copies and, so, that she was doing so and doing so intentionally.

In the first more fleshed out version of the case, I say that the agent who succeeds and who checks ends up having speculative knowledge of marking the 10th copy and has practical knowledge of making 10 copies. Her speculative knowledge is of a type-(iii) fact that potentiates further action, whereas her practical knowledge is of the actualization of the corresponding type-(ii) fact. Her speculative knowledge and her practical knowledge have the same object, namely, the making of the 10 copies. Yet, again, the role of speculative knowledge here is not that of constituting or replacing her practical knowledge but, rather, that of facilitating further action, say, going on and delivering the copies to her boss or just getting on with her day. And these latter actions are things she knows practically by their being grounded in the exercise of her cogent practical cognition. I take this to be shown by the fact that, other things being equal and assuming no change of mind occurs, had she failed and seen this when she checked, she would have gone on to try to mark the rest of the copies using less (and less) interlaced sheets. That is, she will have shifted from the first more fleshed out case to the second. It is this counterfactual shift to the second elaborated case that shows that the first one is no counterexample to the factualist understanding of practical knowledge: her cogent practical cognizing and the support it gets from perception allow her to address potential failure such that she knows "that it will be no accident if [she] ends up having intentionally done what [she] is doing intentionally (or intends to do)" (Small 2012: 136). That is, the agent does not merely aspire to make 10 carbon copies in this case because she is in possession of the relevant type-(ii) facts (read: means), some of which she actively synthesizes and some of which she immediately instantiates. And this is evinced by the counterfactual in which she addresses the relevant *prima facie* failures through additional practical cognizing.

What about a carbon-copier who lacks this procedural knowledge-how to make 10 carbon copies? It is here that aspiration can help us. In such a case, the agent merely aspires to make 10 carbon copies and her aspiration will drive her to calculate means to learning the relevant type-(ii) facts: she will learn how to interlace the carbon-paper with the would-be copies, how to press sufficiently hard on the top sheet, etc. Again, other things being equal and assuming no change of mind occurs, her robust practical cognizing will allow her to avoid or address *prima facie* failures and her ultimate success at making 10 carbon copies will be something she has and exercises practical knowledge of and, so, is doing and doing intentionally.

What about a case in which there is no (more) opportunity for correction? In such a case, we find out more about the agent's procedural knowledge-how: either she knows how to make 10 carbon copies or she knows how to make 10 carbon copies all at once on the first try. The bit of procedural knowledge-how appropriate for such a case is the latter. If the agent has this procedural knowledge-how then what she does is intentional by virtue of her putting this knowledge to work in practical-cognition-*cum*-practical-knowledge, where this manifests as the performance of a skillful action. If the agent only has the former bit of procedural knowledge-how then she merely aspires to make 10 carbon copies all at once on the first try. Her success in such a case is consistent with the factualist understanding of practical knowledge. For her act of making 10 carbon copies all at once on the first try is not intentional. What she does intentionally that culminates in this non-intentional action are her interlacing the sheets, pressing down as hard as she can on the top sheet, etc. Granted, her non-intentional action is constituted by her intentional actions. What makes the difference between her intentional and non-intentional actions is that she only has practical knowledge of the relevant event in its more immediate descriptions. The act of making 10 carbon copies all at once on the first try is "beyond the break" because it is not a matter of course for her that she will succeed (Anscombe 1957/2000: §34). As it stands, this case is still under described: the agent might be practicing with a view to acquiring the skill to make 10 carbon copies all at once on the first try, she might be making a foolhardy attempt for fun, she might have a gun to her head, etc. In any of these cases, so long as the action does not itself result in her acquiring the relevant skill on that occasion, her success is not intentional but a fluke.[139]

Similar considerations apply in some cases where the agent does not check whether she succeeds and where her success is not a matter of course. In such a case, we again find out more about the agent's procedural knowledge-how. If the agent only has procedural knowledge how to make 10 carbon copies then she merely aspires to make 10 carbon copies all at once on the first try and, again, her lucky success is consistent with the factualist understanding of practical knowledge. What is left is the case where the agent has procedural knowledge how to make 10 carbon copies all at once on the first try but where she does not check whether she has succeeded and where success is not a matter of course. Since, *ex hypothesi*, success is not a matter of course in this case, it will not do to say that she simply lacks the skill or is not sufficiently skilled, at least where this means that her having that skill would make success a matter of course. Successful action with the requisite procedural

[139] I take this case to be of a kind with Setiya's fist-clenching case and the dart-throwing case that I consider in the previous Appendix.

knowledge-how but where success is still not a given and where the agent does not check (or where checking does not matter) is the last possible version of the carbon copier counterexample.

Indeed, Shepherd and Carter (2023) proliferate cases with this structure using professional athletes. They take them to undermine the factualist understanding of practical knowledge. Here is their summary of what they call "the initial argument":

> Once one focuses on activities where an agent's success rate grades out somewhere in the middle percentiles, examples of intentional action without [practical] knowledge proliferate. The chief reason is that intentional action and [practical] knowledge have different levels of permissiveness regarding failure in similar circumstances. Intentional action frequently occurs in the face of nearby failure; not so regarding [practical] knowledge. So, the initial argument concludes, [practical] knowledge of *A*-ing is not necessary for intentional *A*-ing (562).

For example, hitting a baseball such that it evades catchers and the like long enough for the hitter to make it to at least first base is not a matter of course even for the best pro hitters. This is called a "base hit." Albert Pujols, the player that Shepherd and Carter first introduce, makes a base hit a little over a third of the time he is at the plate. Moreover, it serves hitters not to check that they've made a good enough hit, since doing so would waste precious time that they could use running.

I contend that the agent in such cases has practical knowledge of *going* to make 10 carbon copies, of *going* to make a base hit, etc. and that it is this practical knowledge that constitutes their intentional action. That is, I deny that "intentional action and [practical] knowledge have different levels of permissiveness regarding failure in similar circumstances" by denying that Pujols can intend to make a base hit (562). Rather, because success is not a matter of course for him (or anyone) he can at most truthfully say, "I am *going to* make a base hit." While this might strike some as odd, there are two things to consider. First, this commits me to the claim that, in cases where Pujols makes a base hit, his doing so is no accident and that he knows that his doing so is no accident despite the fact that his doing so is admittedly chancy. Second, what is lacking on the agent's part in such cases that would in principle allow her to make success a matter of course is speculative knowledge of type-(iii) facts that would aid her corresponding practical knowledge by allowing her to increase the chances of success when she calculates in light of such speculative knowledge.

I take it that the chanciness of a base hit does not thereby make a successful hit an accident. After all, following Shepherd and Carter, Pujols exercises a great deal of "control" in hitting with a view to making a base hit. What makes the hit chancy are features of the game: it is designed such that these hits are chancy. Otherwise, the game would not be entertaining or exhibit the virtues of a competitive sport. And, in agreement with Shepherd and Carter, they are "a bit too chancy to allow for [speculative] knowledge" (574).

However, what I take this chanciness to imply is not that Pujols fails to have *practical* knowledge of his going on to make a base hit when he in fact does so. What it implies is that Pujols does not have *speculative* knowledge of the relevant type-(iii) facts either in cases of success or in those of failure. These facts can be especially fine-grained. And in cases where Pujols fails, had he known the relevant type-(iii) facts and were it that could he adjust his practical cognition accordingly in light of them, he would thereby decrease the chanciness of making a base hit.[140] He could then properly intend to make a base hit. That said, none of this impugns Pujols's cogent practical cognizing of his going to make a base hit achieving the status of an exercise of practical knowledge in cases where he does make the hit despite there being only a 37% chance of his doing so. That is, the Pujols case is unlike the previous carbon copier case, where the agent does not check while lacking the required procedural knowledge-how, because Pujols is sufficiently skilled whereas the carbon copier is not. This is a difference that makes a difference to what we can attribute to them. The latter can only aspire to make 10 carbon copies all at once on the first try in large part because she lacks the relevant skill. Her success would thus be a fluke and, so, non-intentional. By contrast, Pujols's skill is such that his success is not a fluke. It is fully non-accidental despite being admittedly chancy. Yet, because success is still not a matter of course and because he does not check (nor has an opportunity to correct), making the base hit goes "beyond the break." Hence why he can at most truthfully say, "I am going to make a base hit."

So, what of Sheperd and Carter? Their understanding of practical knowledge strikes me as strongly cognitivist. Strong cognitivism about intention and practical knowledge is the view that intentions are a kind of belief and that practical knowledge is a true intention so characterized (see Velleman 1989, Schwenkler 2019, Marušić and Schwenkler 2018). As Shephard and Carter say "'practical knowledge' [...] is arguably just a special case of propositional knowledge" (2023: 568; fn.21). And by this I take it that that they think practical knowledge is an agent's being in a specific (occurrent) mental state, one that is her having a truth-conditional propositional attitude towards the proposition, e.g., "I am doing *A* (vis-à-vis *this*) in order to do *D*," or "I am doing *D* by doing *A* (vis-à-vis *this*)," in cases where she is doing *A*. This cognitive attitude exists separately from the doing, structures it in light of what the proposition conveys about the structure of the desired state of affairs, and is true in virtue of the proposition's being true rather than the agent's making it true. Having assimilated practical knowledge with propositional knowledge, Shepherd and Carter go on to argue that while

[140] I would speculate that in such a case Pujols would be accused of cheating. Even if he had achieved such fine-grained knowledge "legitimately," surely his doing so would give rise to philosophical investigations into the nature of fairness in competitive sports. Hence the claim that the game is designed such that the successful hits of the relevant sorts are chancy.

"intentional action is permissive of risk [of failure]," the corresponding propositional knowledge is subject to standards of safety that prevent it from being so permissive (574).

They suggest that, instead of locating the essence of agency in such truth-conditional propositional attitudes, a proper view of agency locates both its essence and its tolerance for the risk of failure in control, where control is:

> a special kind of mental structure, a structure of specific sorts of attitudes [...] as well as a layered cognitive architecture [...], and in addition a range of tricks and techniques (not only the formation of habits and 'motor schemata' but the capacity to abut one's fallibility by strategic planning [...], including strategic back-up planning [...]), [that] raise our chances of success in an unfriendly world (579).

The idea is to make a shift within a reductive version of the causalist picture in response to the failure of the cognitivist account of practical knowledge to uphold the factualist understanding of such knowledge. I think this move is unnecessary given a denial of cognitivism and causalism. Once we understand practical cognition properly, i.e., as itself a special kind of mental and bodily structure[141] that is constituted by organized clusters of dispositions for actions concerned with identifying and implementing calculative structure, we do not need to decide between "control" as Shepherd and Carter are understanding it here and practical knowledge (however, see chapter 4 of this dissertation). One's successful exercise of "control" over one's actions, in their sense of control, just is her exercising her practical knowledge. As I have been at pains to show, cogent practical cognizing (with appropriate support from speculative knowledge) *includes* the agent's doing what she can to abut her fallibility. Cogent practical cognizing is cogent in part by fixing on ends attainable by the agent and means that she can implement, where the relevant immediately practicable means fall within her repertoire of habits. Likewise, it is cogent and robust in its activity by virtue of admitting of adjustment and refinement of her calculation as she acts. It is only by taking cogent practical cognizing to be some (occurrent) mental state somehow entailing, containing, or fashioned after belief, which exists separately but persists alongside action, that there is any temptation to move away from practical-cognition-*cum*-practical-knowledge as being of the essence of intentional action.

I take this response to commit me to the claim that practical knowledge does not require safety in all of its instances. Knowing that you are doing *A* and why such that you could not easily have been wrong in a similar case is surely preferable to unsafe practical knowledge. And I think safe practical knowledge is paradigmatic

[141] Either or both of which is plausibly (partly) realized in part in the activities of the relevant parts of the neurocognitive architecture.

given how robust practical cognizing is in identifying and rectifying *prima facie* failures. But, as I am sure Shepherd and Carter would admit, there can be non-accidental but chancy success in intentional action. In contexts of sports played at the pro-level, the game and incentives are structured around making success chancy but non-accidental. A pro athlete's skill is most fully on display in such contexts. The appropriate response at this stage is to point out that we have an equivalent for speculative knowledge, namely, quiz shows. It can happen on a show like *Jeopardy!* that a question requires a contestant to do some quick reasoning in order to figure out the answer. That is, the contestant has not memorized every answer that is likely to show up and may not be drawing on previous knowledge to answer the relevant question. In such a case, the pressure of arriving at the answer in time might preclude her doing whatever epistemic work is required to arrive at a corresponding safe belief. But her coming to the relevant unsafe belief and its being true and the correct answer are not an accident. I would be willing to ascribe to the contestant speculative knowledge of the relevant answer on the grounds that, outside of the context of the quiz show, she would have taken measures to found her belief on evidence in a way that made it safe. This would be a kind of epistemic due diligence. Likewise, Pujols, outside of the context of a game of baseball and when practicing, does everything he can to maximize the reliability of making at least a base hit. My overall sense is that worries about safety in particular cases stem from a failure to recognize the larger context of epistemic discursive praxes that scaffold and support our ordinarily coming to safe beliefs. Once that larger context is in view, we can see that cases like Pujols's are not (as I take Shepherd and Carter to be implying) the practical equivalent of skeptical scenarios. Rather, cases like Pujols's are the practical equivalent of pushing yourself to your epistemic limit (say, by being on a quiz show).

Appendix E: The Simple View and Bratman's Video Game

In this Appendix, I'd like to consider one last kind of case, namely, Bratman's (1984) video game case against what he calls the "Simple View". This is a case of unwitting success in which an agent means to do something and succeeds despite having no good reason to think that she would. On the Simple View, if an agent intentionally does *A* then she intends to do *A*. The idea is that there is a common factor between intending to *A* and intentionally doing *A*, namely, the intention to do *A*. Against the Simple View, Bratman devises a counterexample involving an agent playing two connected video games in tandem (381-383). Summarizing, each game is the same: there is a "missile" that is launched upon starting and which the player guides towards a target. The game is so difficult that the agent doubts her ability to succeed at hitting the target. She is ambidextrous and can play both games simultaneously. And here's the twist: the games are connected such that they run simultaneously and if the agent guides both missiles successfully and hits both targets simultaneously, she loses! So, to win the meta-game, the agent has to win one and only one of the two connected games. Yet, because winning each game is difficult to the point of being doubtful regarding her success at winning either, her best strategy is to play both games simultaneously.

Bratman points out that the Simple View cannot maintain that, in this case, the agent simultaneously intends to win the meta-game, intends to win the first game, and intends to win the second game without giving up what he calls "strong consistency" (381). Intentions are strongly consistent relative to beliefs if those intentions "could be put together into an overall plan that is consistent with those beliefs" (380). Strong consistency is a condition on cogent practical cognition: separated from belief, the agent's practical cognizing is liable to be met with one of the kinds of *prima facie* failures and the agent is unlikely to correct. A lack of strong consistency would result in practical cognizing's being a kind of "frictionless spinning in a void," though not the kind that is at issue in McDowell's (1994: 11). We should want to maintain strong consistency among our intentions, then, as cogent practical cognition apt to achieve the status of an exercise of practical knowledge will rely on our having the relevant true beliefs and on our intentions being such that the calculations we make using them are consistent with those beliefs. Since, in the counterexample, the best strategy is to play both games simultaneously, the agent's doing so as a means to winning the meta-game implies on the Simple View that she intends to win the meta-game, to win the first game, and to win the second game. Yet, the conjoined intentions to win the first and second games violate strong consistency with respect to the intention to win the meta-game given the agent's speculative knowledge concerning the nature of their connection. And since any legitimate

account of intention should respect strong consistency, so the argument goes, the counterexample shows that the Simple View is false.

I want to defend the Simple View and I believe the notion of aspiration gives it a way of responding to the supposed counterexample. I want to defend the view because the powers framework adopts a version of it. Indeed, Small (2012), in arguing that the temporal structure of intentional actions in prospect depends on the calculative structure definitional of intentional action in general, says the following:

> [W]hen the concept *do B* is an intentional action concept the legitimacy of its present progressive predication *does* depend on the presence of intention in the agent. If an agent is doing *B* intentionally, then he is going to do *B*, where this consists in his intending to do *B*, of which intention his doing *B* intentionally is the execution (in progress) (189-190; original emphasis).

The agent's intending to do *B* is constitutive of the agent's going to do *B* or having done *B* in this case because the intention is that in virtue of which the agent is presently doing *B*. The idea is that the possibility of intentional action in prospect, e.g., of going to pick up the kids from school, is secured by the calculative structure of intentional action generally. It is secured by this calculative structure because, even in cases where one is now doing the action that is one's end, the goal of completing the action is present and has its sense only in the agent's having a corresponding intention to perform the action to completion. Once that intention is satisfied or given up, the calculative structure shaping the agent's action dissipates. On the powers framework, the imposition of that structure depends on the agent's practically cognizing what she is doing. All that action in prospect adds to this structure is the answer to a *temporal* distance between the agent and her goal. The way that the agent closes that distance is the same as the way that she closes spatial and conceptual distances: cogent practical cognizing. And this means that the agent intends at least some of the means that she takes to closing *any type of distance*, including the immediately practicable means that she implements as an exercise of habit. In intentional actions in progress and in those in prospect, the agent knows that she is doing or going to do that thing, respectively, and that it will be no accident if she succeeds. And it is in virtue of the exercise of such practical knowledge that she is doing or going to do those things and does them or goes on to do them intentionally.

So, there is a sense in which the powers framework adopts the Simple View.[142] It is not in virtue of positing (occurrent) mental states with truth-conditional contents separable from the corresponding action and

[142] Indeed, as Small says in a footnote immediately after the previous quote: "It follows from these remarks that if *S* is doing *B* intentionally, then *S* intends to do *B*, which would presumably count as a version of [...] "the Simple View" [...] [T]he

which exist in the head as, say, a plan. Rather, it is in virtue of the calculative structure of practical cognizing and what that implies about the agent's (practical) understanding of what she is doing that she ends up intending at least some of the means that she takes to her end. This leads us back to the counterexample: if it stands then the Simple View is false and the powers framework is out.

I said that aspiration gives us a way of responding to the supposed counterexample. It does so by allowing us to say that, in the case so described, the agent merely aspires to win the meta-game. She might well also only aspire to win each of the other games and might only intend to play them simultaneously. If this is right then it means that multiple aspirations can relate to each other. In this case, the relationship between the aspiration to win the meta-game and the aspiration to win each of the other games is one of constitution, since winning one and only one of the games just is winning the meta-game. There might well be other relations that obtain between aspirations (and intentions). For instance, an agent might aspire to do two things that relate to each other as cause and effect, where the agent knows *that* the one thing they aspire to do is partly an effect of the other thing they aspire to do. Bratman's player might aspire to win a much more complicated meta-game that involves winning one and only one of the connected games as a first step, where at least some of the other steps are things the agent does not know how to do. Or an agent might also aspire to do something that is a means to something that she intends to do. This kind of case comes up with actions whose means are liable to change where at least sometimes the relevant means are extremely hard to pull off. Procedurally generated levels in a video game serve as a good example. Suppose an agent is playing a game where the levels are procedurally generated. An agent might intend to reach the end of a level insofar as she intends to do this by traversing the level. Yet, she might only be able to aspire to traverse *this iteration* of the level because the procedural generation algorithm designed an iteration of the level that just happens to be so difficult that the agent doubts that she can traverse it. How can she still intend to reach the end of the level? Simple. Say she fails on this attempt because of the difficulty of this iteration. She reaches a "Game Over" screen and selects the option to continue, causing the creation of a new iteration of the level that might well be easy enough for her to intend to traverse it. Because the means that the agent can take are not exhausted by what she aspires to do, she can intend to do something to which what she aspires to do is on a particular occasion a sufficient means.

Returning to Bratman's video game player, the gap in procedural knowledge-how that results in the agent's not knowing that she can win the meta-game (or either of the other games) concerns an *aspect* of the

account of the calculative and temporal structure of intentional action (in progress and in prospect) that I have been developing here constitutes [...] the deep ground for a defensible version of the Simple View" (190: fn116).

action constituting winning either of the connected games, namely, how she can do so *reliably*. If the agent had procedural knowledge-how of this type-(ii) fact then such procedural knowledge-how would effectively close the gap between her playing the connected games and her winning the meta-game, obviating the need to play both games simultaneously. She could then intend both to win the meta-game by winning one of the other games and to win one of the other games by playing it without threat of violating strong consistency. As it stands, then, mere aspirations need not be strongly consistent. In which case, she only intends to play the connected games simultaneously and intending this is strongly consistent. That said, aspirations might be weakly consistent, where this means that they can be put together with the relevant intentions to form an overall plan. Even so, this is consistent with the powers framework's version of the Simple View. The video game player has a coherent plan in the case under consideration given the difficulty of the game: her calculation is made in light of speculative knowledge of the relevant probabilistic facts and she is not aspiring to win *both* of the connected games despite it being true both that she aspires to win the one game and that she aspires to win the other. I contend that the scope of aspiration leaves the conjunction out. One can aspire to do many conflicting things. Practical coherence only requires that aspirations lead to strongly consistent intentions. So, at this point, the player might rightly admit that she does not *really* know what is going to happen and that she is trying to win the meta-game the best way she knows how. This reflects her backing off from assenting to the claims that she will win the meta-game by winning one and only one of the other games and that she will win one and only one of those games by playing both simultaneously. It is plausible that the pressure to back off from such assent stems from recognition that having the corresponding intentions or conjoining her aspirations and relating to them as intentions would amount to a lapse in practical coherence. In any case, the powers framework's version of the Simple View avoids Bratman's counterexample.

Finally, it is worth noting how Bratman's case compares to Shepherd's and Carter's Pujols-type cases. In the latter, Pujols is intending to go make a base hit. He knows how he can do this within the strictures of the game and might have speculative knowledge that he is successful in making a base hit ~37% of the time. In such a case, he cannot intend to make a base hit every time he steps up to the plate nor to make a base hit (flawlessly) without violating strong consistency. Again, it would be more accurate to say that he thinks he is going to make a base hit. This places the action "beyond the break" where the action in "in its more immediate descriptions" include his doing the things he has trained and practiced doing to achieve a base hit rate of ~37%. And this is because success at making a base hit is not a matter of course. "Making a base hit" is not something he can intend just like "winning one and only one of the connected games" and "winning the meta-game" are not

things that the player of Bratman's video game can intend. Pujols, then, lacks both speculative knowledge of whatever type-(iii) facts are relevant to his making a base hit reliably (well above chance) and procedural knowledge-how of the corresponding type-(ii) facts whose exercise in his practical cognizing would just be his making such a hit reliably. It is the same for the video game player with respect to hitting each target reliably. The difference is that, while Pujols can be confident that he is going to make a base hit (despite the chanciness of his success), we are supposing for the sake of argument that the video game player is not confident either that she will win either of the connected games or that she is going to win one and only one of those games. Hence its being the best strategy for her to play both simultaneously. If Pujols lacked confidence, he would make similar strategic moves and look for additional ways of maximizing his chances of making a base hit within the strictures of the game. If one is convinced that safety is a requirement on possessing or exercising knowledge of any kind then they can accept, with Schwenkler (2019), that the knowledge-thesis or cognitive condition applies paradigmatically but does not set out a strict requirement on some action's counting as intentional. If one is convinced that safety is a requirement for possessing speculative knowledge but not a requirement for possessing or exercising practical knowledge then they can accept what I have said about practical knowledge and deny that unsafe but non-accidentally true beliefs count as speculative knowledge.

Appendix F: A Short Critique of the Main Reductive Causalist Accounts of Mental Action

While I did not have space in the main text to consider reductive causalist accounts of mental action, I'd like to do so here. One reason is to give such accounts due space in the dissertation, albeit not in the main text. The other reason is that the issues that I want to argue that the issues that I claimed in Sections 2.4.2 and 4.5.2 plagued causalism apply to reductive causalist accounts of mental action. So, in this appendix, I consider two kinds of account that I think together account for what most reductive causalists think the underlying metaphysics of mental action consists in.[143] According to these accounts, mental actions depend for their performance on attentional or metacognitive mechanisms to (non-deviantly) cause and guide the mental processes constitutive of their performance. I call these accounts "Attention-First Accounts" and "Metacognitive Accounts," respectively. I consider them in turn.

F.1 Attention-First Accounts

F.1.1 Levy's Attention-as-Determinable Account

I focus on the Attention-First Accounts of Wu (2023a,b) and Levy (2023). According to both, because attention to the object of any mental action is at least partly constitutive of performing that action, the agent is selectively directed towards the object of her mental action during performance. Starting with Levy (2023), all mental actions are determinates of which attention is the determinable (91-92). Call this the Attention-as-Determinable Account. The account is meant to explain what Levy calls *Entailment*, where "*A*" stands for an agent, "*V*" for an action, and "*O*" for an object:

[143] Of course, there are also agent-causalist accounts of mental action that I do not here consider. Such accounts, I think, inherit the problems of causalism more generally and are not, to my knowledge, currently in vogue. Notable exceptions are Michael Brent and Thomas Crowther. The latter I do not take particular issue with, since his work provides what is to my mind a suitable explication of listening once we have dropped the background agent-causalist metaphysics. Brent's (2023) account strikes me as problematic for two reasons. First, his objection to Strawson's dilemma is what I identified as the first response in Section 2.4.1. It strikes me as a friendly emendation rather than a genuine objection. Second, he claims that mental agency depends on a sense of agency, understood as the agent's consciously exerting effort. But, as he himself admits, this cannot be a full account: the experience that one has of effortfully doing *A* in doing *A* is in the good case a part of the effortful exercise of agency and, in human beings, functions as part of cogent practical cognizing to promote further calculation and the implementation of means. Intentional action depends on cogent practical cognition. But such cognition is not exhausted by the experience one has of effortfully doing *A* in doing *A*. For such an experience on its own does not account for the calculative structure of intentional action. Moreover, the account cannot make sense of effortlessly doing *A* where the value of "*A*" is a mental action.

Entailment: For a mental act, V, '*A V*-s [preposition] *O*' entails '*A* attends to *O*' (88).

For instance, one's listening to a conversation entails that one attends to that conversation. Similarly, one's calculating the sum of two numbers in a bit of mental arithmetic entails that one attends to those numbers. And so on. Levy takes *Entailment* to be an *explanandum* for any account of attention. Levy's preferred *explanans* is that attention is there whenever one performs a mental action for the same reason that redness is present whenever any determinate shade of red is present. Redness is present whenever scarlet is present because something's being scarlet is a way of its being red. Similarly, every mental action is attentive because every mental action is a way of attending. In the case of mental arithmetic, the agent attends to the application of the relevant equation and their arguments, drawing the solution thereby. The same goes for all mental actions. For instance, in deliberating whether to submit a manuscript to *Philosophical Quarterly*, the agent attends to reasons for and against submitting to that journal and settles on a decision thereby. It follows on Levy's account that mental actions are never inattentive. By contrast, inattentive non-mental action abounds. What I called rote routines in Section 1.2.4 is a case in point: one's inattentive driving one's normal route home is just as much one's action as one's fully attentive driving. Mental actions, on the other hand, "seem precisely to constitute specific ways or modes of paying attention" (91).

Speaking strictly, Levy's account of mental action and his account of attention are non-reductive. So, his account of mental action is not a form of reductive causalism. Nonetheless, I lump him in with Attention-First Accounts because I think that the problems that his account faces are instructive for understanding the potential difficulties for reductive causalism. There are at least three problems with Levy's account. First, only mental actions that comprise the maintenance of some state can plausibly fit the account. Mental actions that comprise the formation or modification of mental states cannot just be ways of attending to the contents of those states. For instance, in reconsidering a belief that *p*, the agent does not merely attend to the reasons or evidence which speak to the likelihood of the truth or falsity of *p*. For it is possible for her to attend to her belief and the relevant reasons and remain, as it were, inferentially inert. To figure out whether *p*, she has to take her reasons to bear on the truth or falsity of *p* and draw the relevant inference thereby. Similarly, in cases of mental arithmetic, merely attending to the operandi and operation is consisting with not doing any such arithmetic. To do mental arithmetic, the agent has to perform the relevant operations. If she did not already know the output then her applying the relevant equation results in the judgment, e.g., that the quantity *r* is the answer. Her

coming to the right judgment in these cases requires her to take an active role in their modification or formation. Attention is insufficient on its own.

The second problem with Levy's account is that it threatens to collapse the distinction between mental action and the occurrence of attentive mental states that are not actions. This problem stems from the commitment that attention is nothing other than the least determinate mental action. Accordingly, attention without mental action is impossible (92-93). Levy is therefore committed to claiming that any content's "popping into one's head" is nonetheless a "non-intentional mental action" (93). He gives as examples of non-intentional mental actions, "talking silently to oneself, gazing idly into the distance, skim-reading billboard signs as one drives past them, and so on" (93). We should add to this list the occurrence of unbidden imagery or memory and merely recognizing things. According to Levy, these mental "movements," though non-intentional, are actions in virtue of involving at least a minimal amount of attention. Otherwise, the agent would not be able to report on having done them or anything about them. Since she can at least sometimes recall having just been gazing idly into the distance, silently speaking to herself, having some unbidden imagery or memory, and so on, she must have been paying some attention. Thus, she must have been performing a non-intentional act of gazing, inner speaking, imagining, or remembering.

However, Levy fails to give a principled distinction between mental actions and mental impulses. Consider skim-reading versus automatic lexical processing or silent speech versus intrusive thoughts which still come in a quasi-auditory "format." For each pair, attention is or can be deployed in both instances but only the former of each counts as a mental action, while the latter is a mental impulse. Recalling Section 2.3.1, the powers framework makes such a distinction by claiming that mental impulses come about in the absence of the agent's at least potentiating(/generating intentional action acorns for) the construction and use of the relevant discursive representations. Both mental impulses and instances of bare recognition are the agent's being merely being primed to perform some mental action. Neither are her performance of any mental action. Returning to the second problem for Levy's account, consider intrusive thoughts. Such thoughts capture attention and can linger despite one's wishing that they would not. Similarly, the pop-out effect in cases of attentional capture suggests that at least some occurrences of attentive mental states are not thereby the performance of any mental action. Levy attempts to avoid this conclusion by claiming that registering content in cases of attentional capture is not yet attentive and therefore rightly not identified with the performance of an action (2023: 93-94). But this move is unmotivated. Without some prior distinction between bare registration and attention proper, we cannot say whether or when automatic lexical processing, intrusive thoughts, unbidden imagery or memory,

and so on are mental actions. The category of non-intentional mental actions is too blunt an instrument. The powers framework provides a preferable alternative. Whether such occurrences count as mental actions depends on whether they are the potentiation of acts of discursive representing sufficient to further potentiate acts of discursive representing that are tied to the same content. If there is no transition, there is no mental action.[144]

This brings us to the third problem with Levy's Attention-as-Determinable Account. Attention of the sort Levy has in mind, namely, as entailing the agent's conscious awareness of acting, is not necessary for the performance of at least some mental actions. Some habitual mental actions, primarily basic mental actions, can be performed in the absence of the agent's conscious awareness and thus in the absence of the relevant sort of attention. Recall our habitual adder from chapter 2. So long as the application of the equation "$n + m$" is a means of judging that "$n + m = r$" that she habitually takes, there is no reason to think that an absence of attention or conscious awareness of it at the time of calculation or judgment discounts her coming to that judgment as her action. So long as comes to recognize the quantity r as the answer from applying the equation, she can claim her applying it as her doing, that is, as something she meant to do or did purposefully. Moreover, her application of the equation can be evaluated against standards of practical rationality and, by being an apt means to her intended end of solving the math problem, meets those standards. The agent's cogent practical cognizing is by stipulation means-end coherent, at least weakly practically consistent with the rest of her mental economy, and is something for which she can in principle cite non-contradictory practical reasons. None of this requires that the agent be consciously aware of applying the equation at the time of calculating for the same reason that such conscious awareness is not required to count her absent-minded driving as an intentional action of hers. All that must be present at the time of performance is the activation of the relevant part(s) of the network of intentional action concepts for applying the equation sufficient to further activate her recognitional skill for discursively representing the quantity r. And assuming that sufficient practice doing (mental) arithmetic has calculatively structured those intentional action concepts into the relevant means and end, there is no reason such cogent practical cognition need be something the agent attends to or is consciously aware of.

[144] In cases of merely maintaining something "before one's mind's eye," such transitions go from potentiation of an act of discursively representing some content to the potentiation of acts of discursively representing that very same content. Covert formation of a content consists in transitions from potentiation of an act of discursively representing some content to the potentiation of acts of discursively representing a distinct content. And covert transformation of a content consists in transitions from potentiation of an act of discursively representing some content to the potentiation of acts of transforming that discursive representation.

The three problems plaguing Levy's account stem from a failure to treat the fundamental idea of action as a constraint on theorizing about mental agency. Recall that, according to that idea, actions are changes that originate or have their source in the agent and that the agent relates to not as an observer or spectator but as an agent (Small 2019: 8*ff.*). The first problem plaguing Levy's account reflects a failure to treat actions as changes that originate or have their source in the agent. If mental actions comprise a category of such changes then they cannot be changes to which the agent merely attends. The agent must be *doing* something in deliberating, inferring, imagining, remembering, doing a bit of mental arithmetic, and so on. The claim that these are merely different ways of attending to the formation, modification, or maintenance of contents misses the agent's differential contributions to such formations, modifications, and maintenance. Were it that she could peer into another agent's mind, her attending to that agent's mental actions should count also as her own mental actions on Levy's account. But this is surely wrong: the second agent is doing things with his mind that the first agent is not. All the first agent is doing is attending. Another way to put the point is that different types of mental action correspond to speciations of the genus of forming, modifying, or maintaining personal-level content-bearing mental states or processes. It is true that they are different ways of doing these three things. But they are not *merely* more determinate ways of doing them. And they are not *merely* ways of attending. For attending as a mental action is a species of maintaining content. Rather, for each most determinable type of mental action, there is a logically independent feature of its performance that is coupled with forming, modifying, or maintaining. I argued that such types roughly correspond to distinct mental faculties: deliberation, inference, imagining, etc. So, Levy is attributing the wrong relation to dispositions for mental actions. They are not *all* related to each other as determinable-determinate. The most general mental skills are related to each other as species of a common genus. And attending is arguably one such species. That said, these species relate to their determinate subskills as determinables.

The second and third problems plaguing Levy's account reflects a failure to treat action as something the agent relates to not as an observer or spectator but as an agent. With respect to the second problem, when asked about what Levy calls her non-intentional mental actions, an agent might well acknowledge that she was doing those things. But she would rightly refuse to give application to Anscombe's "Why" question in explaining what was going on. That is, the agent would acknowledge that she was idly gazing or undergoing intrusive thoughts but would appeal to what Anscombe calls "mental causes" of those happenings and not practical reasons for her doing those things (1957/2000: §§10-15). With respect to the third problem, Levy mislocates the source of the relation that an agent bears to her actions. Conscious awareness is not necessary for

that relation to obtain. After all, an agent relates to her inattentive non-mental actions as an agent. She can say of her inattentive driving that she was taking her normal route home because the workday had ended, because she needed to get home, because she wanted to make dinner, and so on. The same can be said for her inattentive mental actions: she can say of her habitual, basic mental act of summing two numbers that she did it because she wanted to know the answer, because she wanted to tip the waiter for the meal, and so on. The source of the agent's ability to provide practical reasons for the relevant happenings is not her conscious awareness of those happenings but the calculation of means and ends that is at work in the production of those happenings. And such a calculation can be present in the absence of conscious awareness, namely, when the dispositions for doing what she does are already internally calculatively structured and when the agent recognizes some need and some target that makes doing what she does with the latter appropriate for satisfying the former. So, although Levy's account is non-reductive, it nonetheless runs afoul of the fundamental idea of action in much the same way as reductive causalism does.

F.1.2 Wu's Attention-as-Input-Based-Guidance-of-Response Account

I turn now to Wu's account. According to his (2023a) mental actions are "ways of attending to content" (76). To understand what Wu means by this, we have to understand the role of attention in his general account of action. Wu claims here and elsewhere that all actions are the agent's traversing a space defined by constantly updating complexes of perceptual or cognitive inputs, cognitive or motor outputs, and possible input-output mappings (2008, 2011b, 2013b, 2016, 2021, 2023a,b). He calls this "behavior space." Whenever an agent traverses behavior space by implementing some input-output mappings as a non-deviant effect of an intention, her traversing thus is an action. The role of attention in such traversal is the input-mediated guidance of action.[145] For example, a soccer player inundated with sensory and cognitive information might select as input only that information which informs her intentionally kicking the ball. That selection is her attending as a function of her intention to kick. This is attention as a state identical to the input set by her intention. The selected information then informs her response, namely, her moving her body such that she kicks the ball, where her movement is the intended motor output. To use a phrase of James's (1890), the agent's attending is at this point her "mind's taking possession" of information relevant for kicking the ball so as to guide her kicking it (Wu 2023a: 67). This is attending as a process of behavioral guidance. The agent's intention plays the role of

[145] On Wu's account, such guidance, in combination with the appropriate intention, suffices to avoid deviance (2016: 14-17).

directing a state of attention towards some perceptual or cognitive information and biasing the input-output mapping of her kicking as guided by the process of attending. Her kick is an intentional action because the mapping that she in fact implements is a non-deviant effect of, on the one hand, her intention setting her state of attention and biasing that mapping and, on the other, the subsequent attentional guidance of the implementation of that mapping.

For Wu, mental actions are the agent's traversing a behavior space defined by constantly updating complexes of perceptual or cognitive inputs, cognitive outputs, and possible input-output mappings as non-deviant effects of acquiring an intention to so traverse. The role of attention is the same for mental action as it is for non-mental action, namely, as the input-based-guidance of response. Hence the title of this subsection. However, in this case, what attention guides for and what makes the action mental is the mapping's having a (specific) cognitive output rather than a motor output. For example, the agent might intend to figure out whether *p* and has the beliefs that if *q* then *p* and that *q*. The cognitive output is the newly formed belief that *p* that is obtained by implementing the relevant input-output mapping, that is, the inference. Her state of attention towards her beliefs that if *q* then *p* and that *q* transitions into a process of guiding the inference from the conjunction of these beliefs to the conclusion that *p* as a result of acquiring the intention to figure out whether *p*. Since attentional guidance is entailed by intention-based-biasing in at least standard cases of mental action, Wu claims that these actions are "ways of attending to content" and, citing James's (1890) description of attention says:

> to the extent that mental actions involve transitions in states that are intentional where input states inform the output states, we have a type of attentional selectivity to intentional content, the sort of attention that can be described as "taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought" (2023a: 76).

As Wu has it here, he, like Levy, is committed to the claim that there are no inattentive mental actions. On Wu's account, this means that it is never the case that one mentally acts and, yet, there is no input-informed-guidance for the implementation of a specific input-output mapping.

As stated, Wu's account does not fall prey to the problems that plague Levy's account. With respect to the first problem, since all actions involve attentional guidance of input-output mappings on Wu's account, he can allow that what attention guides is the formation and modification of content-bearing personal-level states. Whenever the relevant input-output mapping is one in which the input is used to produce the output or is itself modified, if that mapping is intended then attentional-guidance of its implementation is an intentional mental

action on the account. In contrast to Levy, Wu's reductionism provides him materials to explain how "ways of attending to content" can extend beyond maintaining a state.

With respect to the second problem, Wu provides a principled distinction between mental actions and mental impulses. While attention as both state and process play an essential role in his account of mental agency, intention-based-biasing plays an equally essential role by constraining the scope of intentional agency. As he puts it elsewhere: "we control what we intend" (2013: 254). Traversals of behavior space between perceptual or cognitive inputs and cognitive outputs are mental actions only if they are mappings which the agent can implement as a non-deviant effect of acquiring the relevant intention. Mental impulses involving attentional capture can occur, so long as they are mappings that the agent is not implementing or could not implement as a non-deviant effect of acquiring the relevant intentions. Because Wu does not claim that attending is in every case a mental action, he can allow instances of attentional capture that fall outside of the category of mental impulse.

Finally, with respect to the third problem, because Wu's account is informed by the current neurophysiology and cognitive neuroscience of attention, he has it, in line with the current science, that attention does not entail conscious awareness. When combined with his account of mental action (and action generally), the empirically informed understanding of attention yields the possibility of attention-mediated (habitual) mental action in the absence of conscious awareness at the time of acting. Recall the mental arithmetic case. Fitting it to Wu's apparatus gives us the following: the input is the agent's perceptually registering the numerals "*n*" and "*m*." The output is her judgment that "$n + m = r$," "the answer is *r*," etc. And the input-output mapping is the application of the operation of addition to *n* and *m*. Because attention does not entail conscious awareness, the agent's intention can spotlight the cost of attentional capture which, given the presence of the recurrent opportunity to apply the operation, triggers and then guides the habituated mental calculation routine. Wu's account can allow that attentional capture triggers and guides the implementation of this routine completely outside of the agent's conscious awareness (*cf.* 2023a: 73-74). In that case, the mental arithmetic is what he would call a "passive action." Strictly speaking, contrary to my gloss of habitual mental actions, the calculation in this case would not be intentional on Wu's account. Nonetheless, it appears that his account is in principle consistent with the performance of (habitual) mental actions in the absence of conscious awareness.

The problem with Wu's account is that it depends on a control-based conception of intentional action that, like Levy's account, fails to countenance the fundamental idea of action as a constraint on theorizing. The

problem for Wu is his accounts of agential control and agential guidance. According to the account of agential control, some behavior is agentially controlled when it is a product of what he calls "cognitive integration" between the information or content of an agent's intention as realized by the activity of her executive system and the information encoded in or the content of her (perceptual) taking of the targets of action as realized by the activity of other parts of her neurocognitive architecture. According to the account of agential guidance, some behavior is agentially guided when a certain input guides the agent's response in instantiating an input-output mapping in light of intention biasing that response through cognitive integration. Again, such input, biasing, and integration are all realized by activity spread across the agent's neurocognitive architecture. So, when an agent's intention biases a certain response and the intention is the source of cognitive integration of information or content such that the executive system and whatever other parts of the agent's neurocognitive architecture selectively generates a task-relevant subject-level state, her going on to respond to the target of that state as a result of that integration and in light of attention guiding that response (on the basis of that state) is an intentional action of hers. As he puts it: "intention identifies agentive control" and "[a]ttention is mental guidance in action, the agent's taking things informing response" (2023b: 28, 65).

The problem with this way of understanding agential control, again, is that it fails to treat the agent as the source of the changes that are her actions and fails to make sense of how the agent relates to her action as agent. It fails in the first way for the reasons I outlined in Section 4.5.2. Agential control, on Wu's account, is ultimately a function of the causal efficacy of information or content realized by activity in the executive system. Such activity is supposed to be constitutive of the agent's intention. But, again, merely stating that such information or content so realized is causally efficacious in producing certain responses is not on its own enough to explain what makes such activity or its causal efficacy agential rather than not. And identifying such activity with the agent's intention does not on its own show that it is so identifiable. One can, as Wu does, appeal to agential guidance to fill the gap. But then the question is what difference such guidance consists in that makes the difference between an agent's doing something and something happening to her as mediated by the activity of her neurocognitive architecture. Wu seems to think that the difference consists in its being the agent who attends and, thus, guides the relevant response when she is acting. But Wu understands such attention as being realized or implemented, again, in activity within the agent's neurocognitive architecture, specifically as neural modulations occurring in light the cognitive integration wrought from her intention (2023b: 73*ff.*). Such activity is supposed to be constitutive of the agent's (perceptual) take on things. And this brings us to the question what makes the difference between some activity being the realization or

implementation of the agent's (perceptual) take on things and that activity not being her (perceptual) take on things. The same problems crop up here as did with talk of agential control understood as activity in the executive system. One can say that the difference is that such a take is activity that is generated endogenously or top-down. Or one can say that such activity realizes information or content that we identify with the agent. Or one can say that it is activity that guides responses in an agential way. But none of these will work. The first requires distinguishing endogeneity(/top-down) and exogeneity(/bottom-up) in the activity of the neurocognitive architecture that seems untenable given that activity in the executive system has exogenous(/bottom-up) drivers. The second begs the question against the powers framework and in favor of reductive causalism given that cognitive neuroscientists' use of such representational posits is not (yet) vouchsafed. And the third is a non-starter.

Wu's account fails in the second way for reasons I outlined in Section 2.4.2. That is, the agent simply does not relate to the activity of her neurocognitive architecture as agent. Yet, she does relate as agent to her intentions, to what she is attending to in acting, and to what she is doing in light of her intention and her attending. To say that intention is agential control and is realized by the activity of executive system and that attending is agential guidance and is realized or implemented in the modulations of the activity of the relevant bits of the agent's neurocognitive architecture does not help. For the agent does not relate to her intentions or attendings *as* what Wu suggests realizes or implements those intentions or attendings. And if the relevant information or content-bearing neuronal activities are really meant to be identical to the agent's intending and attending and if the causal efficacy of such activities are really meant to give a specifically *metaphysical* account of agency and action then the agent had better relate to those activities as agent. But, surely, she does not. She does not have unmediated conceptions or thoughts of such neuronal activities such that she can provide practical reasons for their occurrence in light of her calculation what to do. If she could give practical reasons for their occurrence, it would be by way of conceptions or thoughts mediated, say, by her use of a science-fictionesque fMRI or CT scan to track such activities. The identity relations Wu posits, then, get the wrong results.

Despite claiming that there are serious problems for the view, I think much of what I have presented of Wu's account can be salvaged. All that would be required is to forgo the identity claims between the activities in the relevant bits of an agent's neurocognitive architecture and intentions, attendings, and so on. Indeed, Wu claims that agents have "action capacities" consisting in capacities for registering psychological inputs and capacities for responses appropriate to the relevant inputs (2023b: 35). They are complex dispositions whose

manifestation is the implementation of an input-output mapping in the absence of the agent's having to intentionally couple the relevant input and output, say, through deliberation. While I do not think Wu's action capacities line up exactly with my understanding of skills and habits, they are surely closely related notions. And positing such capacities need not come with the identity claims that I think are problematic. Wu can claim that the relevant neuronal activities underwrite but do not realize or implement or constitute the activation or manifestation of action capacities. And Wu can argue that attention is necessarily involved in such activation and manifestation insofar as such activation and manifestation affects the implementation of some input-output mapping rather than others. Attention as selection for action can remain an integral part of intentional action on such a view. So, with what is to my mind some minor retooling, Wu's account can avoid the difficulties that I think currently plague it.

Attention First Accounts of mental action (as they stand) fail to countenance the fundamental idea of action as a constraint on their theorizing. The reason for this failure ultimately stems from the background metaphysics from which they draw their materials and methodology. In Levy's case, while he goes non-reductive, he nonetheless adopts much of his materials and methodology from a fundamentally reductive view of the mind that is not explicitly concerned the agent as the source or originator of her actions nor with her relating to her actions as agent. I think this is borne out in the problems that his view faces. For failure to capture the formation or modification of mental states or processes, failure to distinguish mental action from mental impulse, and failure to capture inattentive mental action altogether point to a treatment of mental states or processes as consisting in states or processes that have truth-conditional contents and that are somehow realized by the activity of the agent's neurocognitive architecture. That he is critical of that treatment elsewhere (2019) does not mean that he has abandoned it entirely. In Wu's case, his going in for a reductive view of mental agency is itself an instance of the way of treating mental states or processes that I find problematic. That said, while I do not think Levy's Attention-as-Determinable Account is salvageable, I think Wu's Attention-as-Input-Based-Guidance-of-Response Account is once the reductive elements of the account are jettisoned.

F.2 Metacognitive Accounts: Proust on Metacognitive Feelings and Metarepresentations

One might claim that the formation, modification, or maintenance of a personal-level content-bearing mental state or process is a mental action only if such formation, modification, or maintenance is (non-deviantly) caused, constituted, or guided by the operation of some metacognitive mechanisms. I take the claim that mental actions depend for their performance on the operation of metacognitive mechanisms to characterize

Metacognitive Accounts. In discussing these accounts, I focus on Proust's (2009, 2013a,b), according to which mental actions are responses to error-signals or bouts of self-doubt that satisfy some epistemic norm.[146] Put simply, mental actions are our controlling our own attitudes (Proust 2009, 2013b: 32). Metacognition comes in at least two stages of mental action. First, the error-signal or self-doubt that motivates a mental action is metacognitive. Second, the agent's evaluation of the response or judging it to have satisfied the relevant epistemic norm is likewise metacognitive. Before connecting metacognition to mental agency, I want to disambiguate metacognition that takes the form of error-signals and evaluation from metacognition that takes the form of self-attribution and self-directed judgment.

The forms that metacognition takes can differ between what has been called, on the one hand, "system-1 metacognition," "procedural metacognition," or "metacognitive feeling" and, on the other, "system-2 metacognition" or "metarepresentation" (for instance, see Proust 2013b; Arango-Muñoz and Bermúdez 2018: 85*ff*.).[147] Metacognitive feelings do not consist in attitudes containing truth-conditional contents that represent the first-order states or processes towards which those feelings are directed. Rather, according to Proust, metacognitive feelings bear a relation of Gricean natural meaning to those first-order states or consist in some intensive gradient-sensitive signal that predicts (degrees of) failure or success with respect to the coming about of the relevant first-order states or with respect to those states' satisfying the relevant epistemic norm(s) (2013b: 22*ff*., 190; see Grice 1957: 377-379 for the distinction between natural and nonnatural meaning). By contrast, metarepresentation does consist in attitudes containing truth-conditional contents that represent lower-order states or processes. Metarepresentations can but need not contain the content of the states they represent. For instance, an agent's metarepresentation can consist in a judgment, say, that her memory about what she had for breakfast is faulty or in a judgment, say, that her memory that she had oatmeal for breakfast is faulty. Either judgment is a form of metarepresentation about the same lower-order state and its content. By contrast, in the midst of attempting to remember what she had for breakfast, the agent can experience various metacognitive feelings that neither represent her memory nor its content. She can experience the so-called "tip-of-the-tongue phenomenon," feel some amount of confidence towards her ability to remember or a lack of confidence, or feel some amount of confidence towards her having forgotten. In the midst of remembering, she might experience her remembering as fluid and easy or as stilted and difficult. After remembering, she might experience some

[146] Proust is not the only proponent of the metacognitive account. But her account is arguably the forerunner. Indeed, several authors appeal to Proust in arguing that metacognition (of some sort) plays an essential role in mental agency (Arango-Muñoz 2013; Arango-Muñoz and Bermúdez 2018; Vierkant 2018, 2022; Mi and Ryan 2020).

[147] Hereafter I use "metacognitive feelings" and "metarepresentation," respectively, except where clarity dictates.

amount of confidence or familiarity towards the memory and its content. And so on. None of these metacognitive feelings are attitudes containing truth-conditional content.

On Proust's (2013a,b) account, the role that metacognition plays in mental agency crosscuts metacognitive feelings and metarepresentations. What matters is that metacognition both motivates and enforces the relevant epistemic norm(s) for the mental action. I consider metarepresentation first. Recall the mental arithmetic case. Suppose this time that the agent judges that she does not know the sum of *n* and *m*, that she knows how to sum, and that she intends to sum *n* and *m*. These metarepresentations of her current mental predicament can motivate her to sum *n* and *m*. She forms the corresponding intention and sets about summing. She can then judge whether she will get it right or, after the fact, that her calculation is (in)correct. The calculation is in this case a mental action partly in virtue of being motivated by a metarepresentation of a lack of knowledge and partly in virtue of being able to satisfy the epistemic norm set by her intention to sum. Were there no motivation, the act of summing either would not have occurred at all or would not have been the non-deviant effect of acquiring the corresponding intention. Were there no way to check for error, there would be no epistemic norm to satisfy, since there would be no way to enforce such a norm (see also Proust 2009).

Turning to metacognitive feelings, they can likewise motivate responses to informational needs and monitor or evaluate those responses with respect to their satisfying the relevant epistemic norm(s). Again, take the mental arithmetic case. Suppose this time that the agent intends to sum *n* and *m* and perceptually registers the numerals "*n*" and "*m*" on a sheet of paper, but the sum does not thereby spring to mind. Suppose further that this failure to form an immediate impulse judgment produces an error-signal. The error-signal motivates the agent to perform the sum. In response, she sets about doing so. Her summing is monitored and evaluated for (in)correctness and either produces a further error-signal or does not, the latter resulting in felt fluency. The summing is in this case a mental action partly in virtue of being motivated by a metacognitive feeling directed towards the informational need and partly in virtue of being able to satisfy the epistemic norm set by her intention. In this case, neither the motivating error-signal, monitoring, nor evaluation are metarepresentational. All three occur as the operation of procedural, sub-personal metacognitive mechanisms directed at the informational need and its (lack of) satisfaction, respectively, and produce corresponding feature-based feelings in the agent. Similar to the case involving metarepresentation, were there no motivation then there would be no

mental action and were there no way to monitor or evaluate the relevant response then there would be no epistemic norm for the response to satisfy.[148]

I take there to be two problems with Proust's metacognitive account. First, the account, like Levy's, fails to allow for inattentive mental action where "inattentive" here means outside of the agent's conscious awareness. Second, the view falls prey to the problems for reductive causalism that I laid out in Sections 2.4.2 and 4.5.2. Starting with the first problem, Proust contends that some mental actions can be performed in the absence of conscious awareness. She writes: "An agent may also only have awareness prior to action rather than when the action occurs, particularly in the case of mental action" (2009: 270). And, in later work, appealing to certain empirical findings:

> Current experimental research suggests, rather, that unconscious information about the dynamics of one's own cognitive system can help one assess the epistemic value of one's own cognitive output, in the domains of perception, memory, and reasoning [...] Neuroscientific evidence shows that cognitive control largely occurs unconsciously [...] Neuroscientific and psychological evidence, however, shows that many epistemic decisions [...] are made unconsciously (2013b: 71-73).

Proust grants that metarepresentation is by its nature conscious, since, after all, the first order mental states or processes that it represents are conscious. Rather, the thought here is that evidence from neuroscience and cognitive neuroscience suggests that mechanisms within system-1 metacognition operate and, thus, exert their influence over mental responses entirely below the threshold of consciousness. In which case, the control over those mental responses in virtue of which they count as mental actions is often exerted without the agent's conscious awareness of having performed the relevant actions at the time of acting. Once we distinguish system-1 metacognition proper from what Proust (2013b: 318) calls the "noetic feelings" which that metacognition produces, we can allow for inattentive mental actions which involve system-1 metacognitive control without producing those feelings. Like Wu's account, then, metacognitive accounts in principle avoid one of the problems plaguing Levy's account.

[148] The dynamics of metacognition can shift throughout the performance of a mental action. The agent can start with a metarepresentation of a task, get engrossed, and shift to metacognitive-feeling-based control. Or she can start with metacognitive feelings and form a metarepresentation of her performance in the midst of acting or after the fact. Assuming performance of the relevant mental action is drawn out for long enough, the agent might well switch back and forth repeatedly between system-1 and system-2 metacognition, trying out a variety of strategies she thinks might provide the answer(s) she is looking for.

Unfortunately, noting that mechanisms within system-1 metacognition often operate below the threshold of consciousness does not suffice to avoid the problem. For, on the one hand, saying of a mental action that it is inattentive does not require claiming that the agent is conscious of her metacognition at the time of acting. And, on the other hand, Proust holds that noetic feelings are functional with respect to motivating and guiding action and are consciously experienced. I consider these points in turn.

When arguing that system-1 metacognition is unconscious, Proust points to experiments in which subjects failed to accurately report what it was that their metacognitive mechanisms were monitoring or evaluating for (2013b: 54-58).[149] When assessing others, subjects failed to recognize that shorter periods of time spent studying an item predicted accuracy of recall for that item because they were unaware that in their own case fluency in processing an item was predictive of successful retention of that item. The thought seems to be that, had they been aware of what fluency was a noetic feeling of, they would have generalized to their assessment of others. Instead, they used an erroneous commonsense heuristic that more time spent studying an item predicts accuracy of recall for that item. However, there is no reason to think that consciousness which results from an intensive gradient-sensitive signaling process entails the ability to accurately report on that process rather than on what the process is directed towards. Pain and bare feelings are cases in point. Both are arguably like system-1 metacognition in that they are non-representational and procedural and are underwritten by intensive gradient-sensitive signals that produce the relevant experiences. Yet, even when pain or, say, inchoate fear is consciously experienced, it does not follow that the pained or afraid can verbally report the quality of their pain or fear or where exactly either is emanating from. Referred pain and fear conditioning show that both can be misdirected and misattributed. This is consistent with a subject's veridically reporting on some occasion exactly what their pain or fear is like and the object(s) towards which either is directed. Sometimes one is fully aware of the quality of one's pains or fear responses. All that the present argument requires is that this is not always the case. The processes that lead to these conscious experiences and what exactly those processes are responses to might be inaccessible to the subject on some occasion. There is no reason to think that the noetic feelings that result from system-1 metacognition would differ from pain or fear in this respect. So, even if the agent is not consciously aware of the metacognitive processes that motivate, monitor, or evaluate her action, still, this does not show that she is not consciously aware of that action. Moreover, even if she is wrong about

[149] One plausible reason for her making this claim is that she claims that system-1 metacognition trades in non-conceptual content, whereas system-2 metacognition trades in conceptual content (2013b, especially chapters 6 and 14). That system-1 metacognition has a nonconceptual representational format means that subjects may not be able to access the heuristics by which system-1 metacognition motivates, monitors, or evaluates performance.

what her noetic feelings are telling her, the agent's having them implies some conscious awareness (however erroneous) of her first-order mental states or processes.

This leads me to the second and more important point, namely, that Proust associates noetic feelings produced by the operation of system-1 metacognition with various dimensions of the agent's awareness of her performance. Granted, Proust expresses some hesitation around saying what the functional role of noetic feelings is with respect to controlling our attitudes and whether those feelings might be had unconsciously (2013b: 58).[150] Yet, despite this hesitation, she claims that noetic feelings are constituents of the agent's emotional experience (166), made available to the agent to inform reasoning (171, footnote 6), and are conscious expressions of neuronal or sub-personal "decisions" (297). Moreover, where there is no delay in response, self-probing, error-signal, or noetic feelings for the agent to experience, Proust claims that the resulting response is a cognitive operation rather than a mental action (193). The reason Proust claims that noetic feelings are consciously experienced yet not truth-conditional and directed at performance is that she takes such feelings to make system-2 metacognition possible and, with it, the propositional attitudes that are supposed to be characteristic of human cognition (301-302). The claim that noetic feelings have this functional role, in combination with the claim that both motivating and evaluative metacognition are necessary steps in every mental action (193), suffices to clinch the entailment from adopting the metacognitive account to denying the possibility of inattentive mental actions. That is, Proust is committed to the claims, first, that the operation of either system-1 or system-2 metacognition prior to and following mental actions is necessary for the relevant first-order states or processes to qualify as mental actions and, second, that the operation of either type of metacognition is sufficient to produce conscious awareness of those actions. She is thus committed to denying the possibility of inattentive mental actions, namely, those that occur in the absence of conscious awareness. And although Proust allows for the possibility of automatic and habitual mental actions, she offers as a response only the claim that her commitment to denying the possibility of inattentive mental action poses no special problem for her account (2013b: 193). I agree. It is a problem her account shares with Levy's and with other forms of reductive causalism more generally.[151]

[150] Proust cites Reder (1996) as a proponent of the view that system-1 metacognition and noetic feelings are both unconscious and that noetic feelings do not have the function of cueing the agent (2013b: 300; see also Reder and Schunn (eds.) 1996). And in other work, Proust discusses routine mental actions (2014: 734-736). Nonetheless, she claims that habitual mental actions, like non-habitual mental actions, depend on the agent's consciously experiencing noetic feelings.
[151] It would be a problem for Brent's account as well had he not restricted its scope.

Finally, Proust's account, like Levy's and Wu's, falls prey to the problems I laid out in Sections 2.4.2 and 4.5.2. For one, Proust ascribes to reductive causalism in her (2001). Second, her identification of system-1 metacognition with neuronal or sub-personal "decisions" suggests that at least mental actions based in metacognitive feelings are supposed to be agential in light of that identification. Both commitments, again, reflect a failure to countenance the fundamental idea of action as a constraint on theorizing about action. The agent does not relate as agent to the activity of the bits of her neurocognitive architecture involved in system-1 metacognition. And claiming that they are neuronal or sub-personal "decisions" in no way shows that the agent herself is the source or originator of the changes to her mental economy constitutive of her metacognitive feeling-based mental actions. If Proust were to abandon her identity claim and were to abandon reductive causalism, her account, like Wu's, could be minimally retooled to avoid these problems. She could claim that metarepresentation and metacognitive feeling are *underwritten* by activity in the relevant bits of the agent's neurocognitive architecture. And she could claim that metarepresentation at least is a function of the activation or manifestation of mental skills for discursive self-representation. I do not take these suggestions as required emendations. Nor do I take those I made for retooling Wu's account as required emendations. I suspect that Proust and Wu could retool their accounts in other ways that nonetheless allow them both to countenance the fundamental idea of action as a constraint on their theorizing. The point is that, as they stand, both accounts run afoul of that idea. And their running afoul of the idea spells trouble for them as accounts of the *metaphysics* of mental action and mental agency.

Bibliography

Aaby, B. and G. Ramsey (2022). The proximate-ultimate distinction and the active role of the organism in evolution. *Biology and Philosophy* 37 (4):1-20.

Addis, D., R. Musicaro, L. Pan, and D. Schacter (2010). Episodic simulation of past and future events in older adults: Evidence from an experimental recombination task. *Psychology and Aging* 25, 369–76.

Addis, D., A. Wong, and D. Schacter (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia* 45, 1363–77.

Agostino, R., J. Sanes, and M. Hallett (1996). Motor skill learning in Parkinson's disease. *Journal of the Neurological Sciences* 139(2): 218-226.

Andrews-Hanna, J. R. (2012). The brain's default network and its adaptive role in internal mentation. *The Neuroscientist* 18(3), 251–270.

Amaya, S. (2013) Slips. *Noûs* 47 (3):559-576.

—— (2020) Out of habit. *Synthese* 198 (12):11161-11185.

Anscombe, GEM. (1957/2000). *Intention*. Cambridge: Harvard University Press.

—— (1999). Practical Truth. *Logos: A Journal of Catholic Thought and Culture* 2 (3):68-76.

—— (2005). Practical Inference. Reprinted in *Human Life, Action and Ethics: Essays by G.E.M. Anscombe*, edited by Mary Geach and Luke Gormally, pp. 109–47. Exeter: Imprint Academic.

Aquinas, *T. Summa theologica,* trans. Fathers of the English Dominican Province (New York: Benziger Brothers, 1911-1925).

Arango-Muñoz, S. (2013). Scaffolded Memory and Metacognitive Feelings. *Review of Philosophy and Psychology* 4 (1):135-152.

Arango-Muñoz, S. and J. P. Bermúdez (2018). Remembering as a mental action. In K. Michaelian, D. Debus, and D. Perrin (Eds.), *New Directions in the Philosophy of Memory*, pp. 75–96. Routledge.

Aristotle & Barnes J. (1984). *The Complete Works of Aristotle: The Revised Oxford Translation*. Princeton University Press.

Aronowitz, S. (forthcoming). Semanticization challenges the episodic-semantic distinction. *British Journal for the Philosophy of Science*.

Badura, C. and H. Wansing (2021) Stit-logic for imagination episodes with voluntary input. *Review of Symbolic Logic*:1-49.

Baier, A. (1970). Act and intent. *Journal of Philosophy* 67 (19):648-658.

Balcerak Jackson, M. (2018). Justification by Imagination. In Fiona Macpherson & Fabian Dorsch (eds.), *Perceptual Imagination and Perceptual Memory*. Oxford University Press. pp. 209-226.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press.

Belnap N., Perloff, M., and Xu, M. (2001) *Facing the Future.* Oxford: Oxford University Press.

Bengson, J. and M. Moffett (eds.) (2011) *Knowing How: Essays on Knowledge, Mind, and Action.* Oxford, England: Oxford University Press USA.

Bermúdez, J.P. and F. Felletti (2021). Introduction: Habitual Action, Automaticity, and Control. *Topoi* 40 (3):587-595.

Bernecker, S. (2010). *Memory: A Philosophical Study*. Oxford University Press.

—— (2015). Visual memory and the bounds of authenticity. In A. Coliva, V. Munz, and D. Moyal-Sharrock (Eds.), *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium*, pp. 445–464. De Gruyter.

Berntsen, D. and A. Bohn (2010). Remembering and forecasting: The relation. *Memory Cognition* 38(3), 265–278.

Berto, F. (2017). Impossible Worlds and the Logic of Imagination. *Erkenntnis* 82 (6):1277-1297.

—— (2018). Aboutness in Imagination. *Philosophical Studies* 175 (8): 1871-1886.

—— (2023). 'Logic Will Get You From A to B, Imagination Will Take You Anywhere' *Noûs* 00. 1-13.

Bickel, L. (2024). Why the performance of habit requires attention. *Mind & Language*, 39(2), 260–270. https://doi.org/10.1111/mila.12467

Blomkvist, A. (2022). Imagination as a skill: A Bayesian proposal. *Synthese* 200 (2): 1-23.

Boorse, C. (1977). Health as a theoretical concept. *Philos Sci* 44(4):542–573. https://doi.org/10.1086/ 288768

Boyle, M. (2011). 'Making up Your Mind' and the Activity of Reason. *Philosophers' Imprint* 11.

—— (2016). Additive Theories of Rationality: A Critique. *European Journal of Philosophy* 24 (3):527-555.

—— (2022). The Rational Role of Perceptual Content. In Matthew Boyle & Evgenia Mylonaki (eds.), *Reason in Nature: New Essays on Themes From John Mcdowell*. Cambridge, Massachusetts: Harvard University Press. pp. 83-110.

Boyle, M. and D. Lavin (2010). Goodness and desire. In Sergio Tenenbaum (ed.), *Desire, Practical Reason, and the Good*. Oxford University Press. pp. 161--201.

Bratman, M. (1984). Two faces of intention. *Philosophical Review* 93 (3):375-405.

Brent, M. (2023). Mental action and the power of effort. In M. Brent and L. Miracchi Titus (eds.) *Mental Action and the Conscious Mind*. New York: Routledge. pp. 100-121.

Briscoe, R. (2011). Mental imagery and the varieties of amodal perception. *Pacific Philosophical Quarterly* 92(2), 153–173.

Briscoe, R. (2018). Superimposed mental imagery: On the uses of make-perceive. In F. Macpherson and F. Dorsch (Eds.), *Perceptual Imagination and Perceptual Memory*, pp. 161– 185.

Brown, D. (2018) Infusing perception with imagination. In F. And Dorsch Macpherson (ed.), *Perceptual Imagination and Perceptual Memory.* pp. 133-160.

Campbel, L. (2017). Two notions of intentional action? Solving a puzzle in Anscombe's Intention. *British Journal for the History of Philosophy* 26 (3): 578-602.

Carruthers, P. (2002). Human creativity: Its cognitive basis, its evolution, and its connections with childhood pretence. *British Journal for the Philosophy of Science* 53 (2): 225-249.

—— (2006). Why Pretend? In Shaun Nichols (ed.), *The Architecture of the Imagination*. Oxford University Press. pp. 89-110.

Carter, J. and T. Poston (2018) *A Critical Introduction to Knowledge-How*. New York: Bloomsbury Academic.

Casas-Roma, J., M. Rodríguez, and A. Huertas (2019) A Common Frame for Formal Imagination. *Minds and Machines* 29 (4):603-634.

Chasid, A. (2017). Imaginative Content, Design-Assumptions and Immersion. *Review of Philosophy and Psychology* 8(2): 259-272.

—— (2019) Imagining in response to fiction: unpacking the infrastructure. *Philosophical Explorations* 23 (1):31-48.

—— (2021a). Not by Imaginings Alone: On How Imaginary Worlds Are Established. *Journal of the American Philosophical Association* 7(2): 195-212.

—— (2021b). Imaginative immersion, regulation, and doxastic mediation. *Synthese* 199 (3-4): 1-43

Chisholm, R. (1964). The descriptive element in the concept of action. *Journal of Philosophy* 61 (20):613-625.

Christoff, K., Z. Irving, K. Fox, R. Spreng, and J. Andrews-Hanna. (2016). Mind-wandering as spontaneous thought: a dynamic framework. *Nature Reviews Neuroscience* 17L 718-731.

Chudnoff, E. (2021). Two Kinds of Cognitive Expertise. *Noûs* 55(2): 270-292.

Clark, A. (2010). Memento's revenge: The extended mind, extended. In Richard Menary (ed.), *The Extended Mind*. MIT Press. pp. 43--66.

Clark, A. and D. Chalmers (1998). The extended mind. *Analysis* 58 (1):7-19.

Comrie, B. (1976). *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*, Cambridge/MA.

Crowther, T. (2009). Perceptual activity and the will. In Lucy O'Brien & Matthew Soteriou (eds.), *Mental actions*. New York: Oxford University Press. pp. 173-191.

Currie, G. (1990). *The Nature of Fiction*. Cambridge University Press.

Currie, G. and I. Ravenscroft. (2002). *Recreative Minds: Imagination in Philosophy and Psychology.* Oxford, GB: Oxford University Press.

Cummins, R. (1975). Functional analysis. *J Philos* 72(November):741–64. https://doi.org/10.2307/2024640

Danto, A. C. (1965). Basic actions. *American Philosophical Quarterly* 2(2), 141–148.

D'Argembeau, A. and M. Van der Linden (2004). Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance. *Consciousness and cognition* 13, 844–58.

—— (2006). Individual differences in the phenomenology of mental time travel: The effect of vivid visual imagery and emotion regulation strategies. *Consciousness and Cognition* 15(2), 342–350.

Davidson, D. (1971). Agent, Action, and Reason. In Robert Williams Binkley, Richard N. Bronaugh & Ausonio Marras (eds.), *Agent, action, and reason*. [Toronto]: University of Toronto Press.

—— (1980/2001). *Essays on Actions and Events*, 2nd Edition. Oxford: Clarendon Press.

De Brigard, F. (2014a). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese* 191(2), 1–31.

—— (2014b). The Nature of Memory Traces. *Philosophy Compass* 9(6): 402-414.

De Brigard, F., D.R. Addis, J.H., Ford, D.L. Schacter, and K.S. Giovanello (2013). Remembering what could have happened: neural correlates of episodic counterfactual thinking. *Neuropsychologia* 51, 2401–2414.

De Brigard, F. and K.S. Giovanello (2012). Influence of outcome valence in the subjective experience of episodic past, future, and counterfactual thinking. *Consciousness and Cognition* 21(3), 1085–1096.

Debus, D. (2008). Experiencing the past: A relational account of recollective memory. *Dialectica* 62(4), 405–432.

—— (2014). Mental time travel: Remembering the past, imagining the future, and the particularity of events. *Review of Philosophy and Psychology* 5(3), 333–350.

—— (2016). XV-shaping our mental lives: On the possibility of mental self-regulation. *Proceedings of the Aristotelian Society* 116(3), 341–365.

Doggett, T. and A. Egan (2007) Wanting things you do not want: The case for an imaginative analogue of desire. *Philosophers' Imprint* 7:1-17.

Dorsch, F. (2009). Judging and the scope of mental agency. In L. O'Brien & M. Soteriou (Ed.), *Mental Actions*. Oxford University Press. pp. 38-71.

—— (2012). *The Unity of Imagining*. De Gruyter.

Douskos, C. (2019). The spontaneousness of skill and the impulsivity of habit. *Synthese* 196 (10):4305-4328.

Drayson, Z. (2014). The Personal/Subpersonal Distinction. *Philosophy Compass* 9 (5):338-346.

Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. MIT Press.

Dreyfus, H. (2006). Overcoming the myth of the mental. *Topoi* 25 (1-2):43-49.

Dunlosky, J., and R.A. Bjork, (2008). *Handbook of metamemory and memory*. Psychology Press.

Dunlosky, J., and S. Tauber, (2016). *The Oxford Handbook of Metamemory*. Oxford University Press.

Evnine, S. (2008). Modal epistemology: Our knowledge of necessity and possibility. *Philosophy Compass* 3(4), 664–684.

—— (2016). *Making Objects and Events: A Hylomorphic Theory of Artifacts, Actions, and Organisms.* Oxford, United Kingdom: Oxford University Press UK.

Falvey, K. (2000) Knowledge in intention. *Philosophical Studies* 99 (1):21-44.

Fantl, J. (2008) Knowing-how and knowing-that. *Philosophy Compass* 3 (3):451–470.

Favela L.H. and E. Machery (2023). Investigating the concept of representation in the neural and psychological sciences. *Front. Psychol*. 14:1165622. doi: 10.3389/fpsyg.2023.1165622Fernández, J. (2019). *Memory: A Self-Referential Account*. New York, NY, USA: Oxford University Press.

Fiocco, M. (2007). Conceivability, Imagination and Modal Knowledge. *Philosophy and Phenomenological Research* 74 (2): 364-380.

Fodor, J. (1975) *The Language of Thought*. Cambridge: Harvard University Press.

Ford, A. (2015). The arithmetic of intention. *American Philosophical Quarterly* 52(2), 129–143.

—— (2016). On What Is in Front of Your Nose. *Philosophical Topics* 44 (1):141-161.

Foultier, A. (2022). Letting the Body Find Its Way: Skills, Expertise, and Bodily Reflection. *Phenomenology and the Cognitive Sciences*:1-22.

Frankfurt, H. (1978). The Problem of Action. *American Philosophical Quarterly* 15: 157-62.

Fridland, E. (2014). They've lost control: reflections on skill. *Synthese* 191 (12):2729-2750.

—— (2015). Knowing-how: Problems and Considerations. *European Journal of Philosophy* 23 (3):703-727.

—— (2017a). Automatically minded. *Synthese* 194(11): 4337–4363.

—— (2017b). Skill and motor control: intelligence all the way down. Philosophical Studies 174 (6):1-22.

—— (2019). Longer, smaller, faster, stronger: On skills and intelligence. *Philosophical Psychology* 32 (5):759-783.

—— (2021). Skill and strategic control. *Synthese* 199 (3-4):5937-5964.

Frost, K. (2016). The Antinomy of Basic Action. In *Time and the Philosophy of Action*, eds. Roman Altshuler and Michael J. Sigrist, 37–51. Routledge.

Galton, A. (1984). *The logic of aspect: an axiomatic approach*. Oxford [Oxfordshire]: Clarendon Press.

—— (2006). On the Process of Coming into Existence. *The Monist* 89 (3): 294-312.

Garson, J. (2017). Against organizational functions. *Philos Sci* 84(5):1093–1103. https://doi.org/10.1086/694009

Gauker, C. (2020) On the Difference Between Realistic and Fantastic Imagining. *Erkenntnis* 87 (4):1563-1582.

—— (2021) Imagination constrained, imagination constructed, *Inquiry,* DOI: 10.1080/0020174X.2021.1933748

Gendler, T. (2000) The Puzzle of Imaginative Resistance. *Journal of Philosophy* 97 (2):55.

—— (2003) On the relation between pretense and belief. In Matthew Kieran & Dominic McIver Lopes (eds.), *Imagination Philosophy and the Arts*. Routledge. pp. 125--141.

—— (2006a) Imaginative resistance revisited. In Shaun Nichols (ed.), *The Architecture of the Imagination*. Oxford University Press. pp. 149-173.

—— (2006b) Imaginative contagion. *Metaphilosophy* 37 (2):183-203.

—— (2010) *Intuition, Imagination, and Philosophical Methodology*. Oxford: Oxford University Press UK.

Giordani, A. (2019). Axiomatizing the Logic of Imagination. *Studia Logica* 107 (4): 639-657.

Godfrey-Smith, P. (1992). Indication and adaptation. *Synthese* 92 (2): 283-312.

Goldman, A. I. (1970). *A Theory of Human Action*. Englewood Cliffs, NJ, USA: Princeton University Press.

—— (2006) *Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading*. Oxford: Oxford University Press.

Goldwasser, S. (2022). Memory as skill. *Review of Philosophy and Psychology*, 1–24.

—— (2023). Standard Aberration: Cancer Biology and the Modeling Account of Normal Function. Biology and Philosophy 38 (1):(4) 1-33.

Gregory, D. (2004) Imagining possibilities. *Philosophy and Phenomenological Research* 69 (2):327–348.

—— (2020) Imagery and Possibility. *Noûs* 54 (4):755-773.

Grice, H. (1957). Meaning. *Philosophical Review* 66 (3):377-388.

Hassabis, D. and E. A. Maguire (2009). The construction system of the brain. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1521), 1263–1271.

Heal, J. (2003). *Mind, Reason and Imagination: Selected Essays in Philosophy of Mind and Language*. Cambridge: Cambridge University Press.

Hegland, K.W., D.C. Bolser, and P.W. Davenport (2012). Volitional control of reflex cough. *Journal of applied physiology* (Bethesda, Md. : 1985), 113(1), 39–46. https://doi.org/10.1152/japplphysiol.01299.2011

Hendrickx, M. (2023). Agentially Controlled Action: Causal, not Counterfactual. *Philosophical Studies* 180 (10-11):3121-3139.

Hess, L. (2006). I would like to play but I do not know how: a case study of pretend play in autism. *Child Language Teaching and Therapy* 22 (1): 97-116.

Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly* 87 (1):45-74.

—— (2009). Believing at Will. *Canadian Journal of Philosophy, Supplementary Volume* 35 (sup1):149-187.

Hoffman, K. L. and B. L. McNaughton (2002). Coordinated reactivation of distributed memory traces in primate neocortex. *Science* 297(5589), 2070–2073.

Hopkins, R. (2014). Episodic memory as representing the past to oneself. *Review of Philosophy and Psychology* 5(3), 313–331.

—— (2018). Imagining the past: On the nature of episodic memory. In F. M. F. Dorsch (Ed.), *Memory and Imagination*, pp. 46–71. Oxford University Press.

—— (2023). Remembering and Imagining as Agential Powers. In A. Berninger and I.V. Ferran (Eds.), *Philosophical Perspectives on Memory and Imagination*, pp. 212-229. New York: Routledge.

—— (2024). *The Profile of Imagining*. Oxford: Oxford University Press.

Hornsby, J. (1980). *Actions.* Routledge and Kegan Paul.

—— (2013). Basic Activity. *Aristotelian Society Supplementary Volume* 87 (1):1-18.

Hume, D. (1739/2007) *A Treatise of Human Nature: a Critical Edition*. Oxford: Clarendon Press.

—— (1777/2000) *An enquiry concerning human understanding: a critical edition.* New York: Oxford University Press.

Ichikawa, J. and B. Jarvis (2012) Rational Imagination and Modal Knowledge. *Noûs* 46 (1):127-158.

Irish, M. and O. Piguet (2013). The pivotal role of semantic memory in remembering the past and imagining the future. *Frontiers in Behavioral Neuroscience* 7.

Irving, Z. (2016). Mind-wandering is unguided attention: accounting for the "purposeful" wanderer. *Philosophical Studies* 173 (20: 547-571.

Irving, Z. and A. Glasser (2019). Mind-wandering: A philosophical guide. *Philosophy Compass* 15 (1): 15:e12644.

Irving, Z., A. Glasser, A. Gopnik, V. Pinter, and C. Sripada. (2020). What Does "Mind-Wandering" Mean to the Folk? An Empirical Investigation. *Cognitive Science* 44 (10): e12908.

James, W. (1980). *The Principles of Psychology*. London, England: Dover Publications.

Jessop, R. C. Horowicz, and L. Dibble (2006). Motor learning and Parkinson disease: Refinement of movement velocity and endpoint excursion in a limits of stability balance task. *Neurorehabilitation and neural repair*, 20(4), 459–467. https://doi.org/10.1177/1545968306287107

Johnson-Laird, P. (1983) *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness.* Cambridge, MA: Harvard University Press.

Kelley, M. (2024). How to perform a nonbasic action. *Noûs* 58 (1).

Kenny, A. (1963). *Action, Emotion And Will*. Ny: Humanities Press.

Kind, A. (2005). Imagery and Imagination. *Internet Encyclopedia of Philosophy*. https://iep.utm.edu/imagery/

—— (2013). The Heterogeneity of the Imagination. *Erkenntnis*, 78(1), 141–59.

—— (2016a). The Snowman's Imagination. *American Philosophical Quarterly,* 53 (4):341-348.

—— (2016b). Imagining under constraints. In Amy Kind & Peter Kung (eds.), *Knowledge Through Imagination*. Oxford University Press. pp. 145-159.

——(2018). How imagination gives rise to knowledge. In F. Macpherson and F. Dorsch (Eds.), *Perceptual Imagination and Perceptual Memory*, pp. 227-246. Oxford University Press.

—— (2019). Mary's Power of Imagination. In Sam Coleman (ed.), *The Knowledge Argument*. Cambridge University Press, pp. 161-179.

—— (2020). The skill of imagination. In E. Fridland and C. Pavese (Eds.), *Routledge Handbook of Skill and Expertise*, pp. 335–346. Routledge.

—— (2022). Learning to imagine. *British Journal of Aesthetics* 62(1), 33–48.

Kind and Badura (eds.) (2021). *Epistemic Uses of Imagination*. New York: Routledge.

Kind, A. and P. Kung (eds.) (2016). *Knowledge Through Imagination*. Oxford, United Kingdom: Oxford University Press UK.

Kirsh, D. and M. Paul (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science* 18 (4):513-49.

Klein, S. B., J. Loftus, and J. F. Kihlstrom (2002). Memory and temporal experience: the effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition* 20(5), 353–379.

Kosslyn, S. (1994). *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, MA: MIT Press.

—— (2005) Mental images and the Brain. *Cognitive Neuropsychology*, 22:3-4, 333-347

Kung, P. (2010). Imagining as a Guide to Possibility. *Philosophy and Phenomenological Research* 81 (3): 620-663.

LaBar, K. S. and R. Cabeza (2006). Cognitive neuroscience of emotional memory. *Nature Reviews Neuroscience* 7(1), 54–64.

Langland-Hassan, P. (2015). Imaginative attitudes. *Philosophy and Phenomenological Research* 90(3), 664–686.

—— (2016) On Choosing What to Imagine. In A. Kind & P. Kung (Eds.), *Knowledge through Imagination* (pp. 61–84). Oxford: Oxford University Press.

—— (2020). *Explaining Imagination*. Oxford: Oxford University Press.

—— (2021). What sort of imagining might remembering be? *Journal of the American Philosophical Association* 7(2), 231–251.

—— (2022a). Propping up the causal theory. *Synthese* 200 (2):1-27.

—— (2022b). Remembering, Imagining, and Memory Traces: Toward a Continuist Causal Theory. In Christopher McCarroll, Kourken Michaelian & Andre Sant'Anna (eds.), *Current Controversies in Philosophy of Memory*. Routledge.

—— (2023). Remembering and imagining: The attitudinal continuity. In A. Berninger and Íngrid Vendrell Ferran (Eds.), *Philosophical Perspectives on Memory and Imagination.* London: Routledge.

Lavin, D. (2013). "Must There Be Basic Action?" *Noûs* 47 (2): 273–301.

Levy, Y. (2013). Intentional action first. *Australasian Journal of Philosophy* 91(4): 705-718

—— (2019). What is 'mental action'? *Philosophical Psychology* 32 (6):971-993.

—— (2023). The Most General Mental Act. In Michael Brent & Lisa Miracchi (eds.), *Mental Action and The Conscious Mind*. Routledge. pp. 79-99.

Liao, S. and T. Gendler (2019). Imagination. *Stanford Encyclopedia of Philosophy*.

Locke, D. (1974). X*—Natural Powers and Human Abilities. *Proceedings of the Aristotelian Society* 74 (1):171-187.

Ludvig, D., I. Cathers, and R.E. Kearney (2007). Voluntary modulation of human stretch reflexes. *Experimental brain research*, 183(2), 201–213. https://doi.org/10.1007/s00221-007-1030-0

Lynch, K. (2019). Knowing how, basic actions, and ways of doing things. *Inquiry: An Interdisciplinary Journal of Philosophy* 62(8), 956–977.

Marcus, E. (2012). *Rational causation*. Cambridge, Mass.: Harvard University Press.

—— (2018). Practical Knowledge as Knowledge of a Normative Judgment. *Manuscrito* (4):319-347.

—— (2019). Reconciling Practical Knowledge with Self-Deception. *Mind* 128 (512):1205-1225.

—— (2021). *Belief, Inference, and the Self-Conscious Mind*. Oxford, United Kingdom: Oxford University Press.

—— (2022a). Anscombe and The Difference Rationality Makes. In Adrian Haddock & Rachael Wiseman (eds.), *The Anscombean Mind*. Routledge.

—— (2022b). Rational Agency. In Luca Ferrero (ed.), *The Routledge Handbook of Philosophy of Agency*. New York, NY: Routledge. pp. 118-124.

Martin, C.B. and M. Deutscher (1966). Remembering. *Philosophical Review* 75 (April):161-96.

Martin, V. C., D. L. Schacter, M. C. Corballis, and D. R. Addis (2011). A role for the hippocampus in encoding simulations of future events. *Proceedings of the National Academy of Sciences* 108(33), 13858–13863.

Marušić, B. and J. Schwenkler (2018). Intending is Believing: A Defense of Strong Cognitivism. *Analytic Philosophy* 59 (3):309-340.

Matsuda, T., T.Y. Hiyama, K. Kobayashi, et al. Distinct CCK-positive SFO neurons are involved in persistent or transient suppression of water intake. *Nat Commun* 11, 5692 (2020). https://doi.org/10.1038/s41467-020-19191-0

Mayr, E. (1961). Darwin's Biological Work: Some Aspects Reconsidered. P. R. Bell. *Isis* 52 (3):433-435.

McCarroll, C. (2018). *Remembering From the Outside: Personal Memory and the Perspectival Mind*. Oxford: Oxford University Press.

McDonough, I. M. and D. A. Gallo (2010). Separating past and future autobiographical events in memory: Evidence for a reality monitoring asymmetry. *Memory Cognition* 38(1), 3–12.

McDowell, J. (1996). *Mind and World*. Cambridge: Harvard University Press.

—— (2008). "Avoiding the Myth of the Given", in: J. Lindgaard (ed.), *John McDowell: Experience, Norm, and Nature*, Oxford.

—— (2010). What is the content of an intention in action? *Ratio* 23(4), 415–432.

—— (2011). Some remarks on intention in action. *The Amherst Lecture in Philosophy* 6, 1–18.

—— (2022). How receptive knowledge relates to practical knowledge. In Adrian Haddock & Rachael Wiseman (eds.), *The Anscombean Mind*. New York, NY, USA: pp. 213-224.

McGinn, C. (2004). *Mindsight: Image, Dream, Meaning.* Cambridge: Harvard University Press.

Mele, A. (2009). Mental action: A case study. In L. O'Brien & M. Soteriou (Ed.), *Mental Actions*. Oxford University Press. pp. 17-37.

Meskin A. and J. Weinberg (2003) Emotions, Fiction, and Cognitive Architecture. *British Journal of Aesthetics*, 43(1), 18–34.

Mi, C. and S. Ryan (2020). Skilful reflection as a master virtue. *Synthese* 197 (6):2295-2308.

Michaelian, K. (2016a). Against discontinuism: Mental time travel and our knowledge of past and future events. In K. Michaelian, S. B. Klein, and K. K. Szpunar (Eds.), *Seeing the Future: Theoretical Perspectives on Future-Oriented Mental Time Travel,* pp. 62–92. Oxford University Press.

—— (2016b). *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. MIT Press.

—— (2021). Imagining the past reliably and unreliably: towards a virtue theory of memory. *Synthese* 199(3-4): 7477-7507

—— (forthcoming). Radicalizing simulationism: Remembering as imagining the (nonpersonal) past. *Philosophical Psychology*, 1–27.

Michaelian, K., D. Perrin, and A. Sant'Anna (2020). Continuities and discontinuities between imagination and memory: The view from philosophy. In A. Abraham (Ed.), *The Cambridge Handbook of Imagination*. Cambridge University Press.

Millikan, R. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. MIT Press.

—— (1995). Pushmi-pullyu representations. *Philosophical Perspectives* 9:185-200.

—— (2006). Styles of Rationality. In Susan Hurley & Matthew Nudds (eds.), *Rational Animals?* Oxford University Press.

Montero, B. (2016). *Thought in Action: Expertise and the Conscious Mind.* Oxford, United Kingdom: Oxford University Press UK.

Moran, R. (2004). Anscombe on 'Practical Knowledge' in J. Hyman and H. Steward (eds.), *Royal Institute of Philosophy Supplement.* Cambridge University Press, pp. 43-68.

Morton, A. (2013). *Emotion and Imagination*. Polity.

Mossio, M. C. Saborido, and A. Moreno (2009). An organizational account of biological functions. *British J Philos Sci* 60(4):813–841. https://doi.org/10.1093/bjps/axp036.

Müller, A. (1979). How Theoretical Is Practical Reason?'. *In Intention and Intentionality: Essays for G. E. M. Anscombe*, edited by Cora Diamond and Jenny Teichman, pp. 91–108. Ithaca, NY: Cornell University Press.

Munro, D. (2021a). Imagining the actual. *Philosophers' Imprint* 21(17).

—— (2021b). Remembering the past and imagining the actual. *Review of Philosophy and Psychology* 12(2).

Munro, D. and M. Strohminger. (2021). Are We Free to Imagine What We Choose? *Synthese* (5-6): 1-18.

Muslimovic, D., B. Post, J.D. Speelman, and B. Schmand. (2007). Motor procedural learning in Parkinson's disease, *Brain*, Volume 130, Issue 11, Pages 2887–2897, https://doi.org/10.1093/brain/awm211

Myers, J. (2021a) The epistemic status of the imagination. *Philosophical Studies* 178(10), 3251–3270.

—— (2021b) Reasoning with imagination. In A. Kind and C. Badura (Eds.), *Epistemic Uses of Imagination*. Routledge

—— (2024). Imagination as a source of empirical justification. *Philosophy Compass* 19 (3):e12969.

—— (forthcoming a). How Imagination Informs. *Philosophical Quarterly*.

—— (forthcoming b). Imaginative Beliefs. *Inquiry: An Interdisciplinary Journal of Philosophy*.

Mylopoulos, M. (2022) Oops! I Did it Again: The Psychology of Everyday Action Slips. *Topics in Cognitive Science* 14 (2):282-294.

Nanay, B. (2010). Perception and imagination: Amodal perception as mental imagery. Philosophical Studies 150(2), 239–254.

—— (2016). Imagination and perception. In A. Kind (Ed.), Routledge Handbook of the Philosophy of Imagination, pp. 124–134. Routledge.

—— (2021). Unconscious Mental Imagery. Philosophical Transactions of the Royal Society B 376 (1817):20190689.

—— (2023). Memory Imagery: Philosophy, Psychology, Neuroscience. Oxford University Press.

Neander, K. (2017). *A Mark of the Mental: A Defence of Informational Teleosemantics*. MIT Press, Cambridge, USA

Nichols, S. (2004) Imagining and Believing: The Promise of a Single Code. *Journal of Aesthetics and Art Criticism*, 62, 129–39.

Nichols, S. and S. Stitch (2000) A Cognitive Theory of Pretense. *Cognition*, 74, 115–47.

—— (2003) *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding of Other Minds*. Oxford: Oxford University Press

Nielsen, M. and T. Christie. (2007). Adult modeling facilitates young children's generation of novel pretend acts. *Infant and Child Development* 17 (2): 151-162.

Noordhof, P. (2002). Imagining objects and imagining experiences. *Mind and Language* 17 (4): 426-455.

—— (2018). Imaginative Content. In Fiona Macpherson & Fabian Dorsch (eds.), *Perceptual Imagination and Perceptual Memory*. Oxford: pp. 96-129.

O'Brien, L. (2017). Actions as Prime. *Royal Institute of Philosophy Supplement* 80:265-285.

Paul, S. (2009). Intention, belief, and wishful thinking: Setiya on "practical knowledge". *Ethics* 119 (3):546-557.

Pavese, C. (2015). Practical Senses. *Philosophers' Imprint* 15.

—— (2020). Practical Representation. In Ellen Fridland & Carlotta Pavese (eds.), *The Routledge Handbook of Philosophy of Skill and Expertise*. New York, NY: Routledge.

—— (2023). Knowledge-how. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Summer 2021 Edition).

Peacocke, A. (2021). Mental action. *Philosophy Compass* 16 (6):e12741.

—— (2023). How to Think Several Thoughts at Once: Content Plurality in Mental Action. In Michael Brent & Lisa Miracchi Titus (eds.), *Mental Action and the Conscious Mind*. New York, NY: Routledge. pp. 31-60.

Peacocke, C. (1985). Imagination, experience, and possibility. In John Foster & Howard Robinson (eds.), *Essays on Berkeley: a tercentennial celebration*. New York: Oxford University Press.

—— (2009). Mental action and self-awareness. In Lucy O'Brien & Matthew Soteriou (eds.), *Mental actions*. New York: Oxford University Press.

Pendt, L.K., I. Reuter, and H. Müller (2011). Motor skill learning, retention, and control deficits in Parkinsons disease. *PLoS ONE* 6:e21669 10.1371/journal.pone.0021669

Peiss, D. P., D. Cosmelli, and J. C. Kaufman (2020). *Creativity and the Wandering Mind: Spontaneous and Controlled Cognition*. Academic Press.

Perrin, D. (2016). Asymmetries in subjective time. In K. Michaelian, S. B. Klein, and K. K. Szpunar (Eds.), *Seeing the future: Theoretical perspectives on future-oriented mental time travel*., pp. 39–61. New York, NY, US: Oxford University Press.

Perrin, D. and K. Michaelian (2017). Memory as mental time travel. In S. Bernecker and K. Michaelian (Eds.), *Routledge Handbook of Philosophy of Memory*, pp. 228–239. Routledge.

Picciuto, E. and P. Carruthers. (2016). The Origins of Creativity. In E.S. Paul and S.B. Kaufmen (eds.) *The Philosophy of Creativity: New Essays.* Oxford University Press, pp. 199-223.

Proust, J. (2001). A plea for mental acts. *Synthese* 129(1), 105–128.

—— (2009). Is there a sense of agency for thought? In Lucy O'Brien & Matthew Soteriou (eds.), *Mental actions*. New York: Oxford University Press.

—— (2013a). Mental acts as natural kinds. In Till Vierkant, Julian Kieverstein & Andy Clark (eds.), *Decomposing the Will*. Oxford University Press. pp. 262-282.

—— (2013b). *The Philosophy of Metacognition: Mental Agency and Self- Awareness*. New York: Oxford University Press.

Proust, M. (2003). *Swann's Way*. New York :Modern Library

Radford, C. and M. Weston. (1975). How Can We Be Moved by the Fate of Anna Karenina? *Proceedings of the Aristotelian Society, Supplementary Volumes* 49: 67-93.

Raichle, M. E. (2015). The brain's default mode network. *Annual Review of Neuroscience* 38(1), 433–447.

Rasmussen, A. S. and D. Berntsen (2013). The reality of the past versus the ideality of the future: emotional valence and functional differences between past and future mental time travel. *Memory Cognition* 41(2), 187–200.

Recanati, F. (2007). *Perspectival Thought: A Plea for Moderate Relativism*. New York: Oxford University Press.

—— (2010). *Truth-Conditional Pragmatics*. Oxford, GB: Oxford University Press.

Reder, L.M. (1996). *Implicit memory and metacognition*. Hillsdale, NJ: Erlbaum

Reder, L.M. and C.D. Shunn (1996). Metacognition does not imply awareness: Strategy choice is governed by implicit learning and memory. '. In L. M. Reder (ed.) *Implicit Memory and Metacognition*, 45–78. Hillsdale, NJ: Erlbaum.

Robins, S. (2016). Representing the past: Memory traces and the causal theory of memory. *Philosophical Studies* 173(11), 2993–3013.

—— (2020). Defending discontinuism, naturally. *Review of Philosophy and Psychology* 11(2), 469–486.

—— (2022). Episodic memory is not for the future. In, A. Sant'Anna, C. McCarroll, & K. Michaelian (Ed.) *Current Controversies in the Philosophy of Memory.* Routledge (pp. 166–184).

Rosenbaum, R. S., S. Köhler, D. Schacter, M. Moscovitch, R. Westmacott, S. Black, F. Gao, and E. Tulving (2005). The case of k.c.: Contributions of a memory-impaired person to memory theory. *Neuropsychologia* 43, 989–1021.

Rothstein, S. (2004). *Structuring events: a study in the semantics of aspect*. Malden: Blackwell.

Ryle, G. (1946/2009). Knowing How and Knowing That: The Presidential Address. *Proceedings of the Aristotelian Society* 46:1 - 16.

—— (1949). *The Concept of Mind*. New York: University of Chicago Press.

Sant'Anna, A. (2023). Is remembering constructive imagining? [Unpublished Manuscript].

Schacter, D. L. and D. R. Addis (2007a). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philos Trans R Soc Lond B Biol Sci* 362(1481), 773–786.

—— (2007b). The ghosts of past and future. *Nature* 445(7123), 27–27.

Schacter, D. L., R. G. Benoit, F. De Brigard, and K. K. Szpunar (2015). Episodic future thinking and episodic counterfactual thinking: Intersections between memory and decisions. *Neurobiology of Learning and Memory* 117, 14–21.

Schneider, W. and R. Shiffrin (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review* 84 (1):1-66.

Schwenkler, J. (2015). Understanding 'Practical Knowledge'. *Philosophers' Imprint* 15 (15): 1-32.

—— (2019). *Anscombe's Intention: A Guide*. New York, USA: Oxford University Press.

Setiya, K. (2007). *Reasons Without Rationalism*. Princeton University Press.

—— (2008). Practical knowledge. *Ethics* 118 (3):388-409.

Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Oxford University Press.

Sellars, W. (1956). Empiricism and the philosophy of mind. *Minnesota Studies in the Philosophy of Science* 1:253-329.

Shea, N. (2018). *Representation in Cognitive Science*. Oxford University Press.

Shepherd, J. (2015). Deciding as Intentional Action: Control over Decisions. *Australasian Journal of Philosophy* 93 (2):335-351.

—— (2019). Why does the mind wander? *Neuroscience of Consciousness*, Volume 2019, Issue 1, niz014.

—— (2021). *The shape of agency: Control, action, skill, knowledge*. Oxford: Oxford University Press.

—— (2023). Disappearing agents, mental action, rational glue. In Michael Brent & Lisa Miracchi Titus (eds.), *Mental action and the conscious mind*. New York: Routledge. pp. 14-30.

Shepherd, J. & C.J. Adam (2023). Knowledge, practical knowledge, and intentional action. *Ergo: An Open Access Journal of Philosophy,* 9:556-583.

Small, W. (2012). "Practical Knowledge and the Structure of Action" In Günter Abel & James Conant (eds.), *Rethinking Epistemology Volume 2*. Berlin, Germany: De Gruyter. pp. 133-227.

—— (2014). The Transmission of Skill. *Philosophical Topics* 42 (1): 85-111.

—— (2017a). Agency and Practical Abilities. *Royal Institute of Philosophy Supplement* 80:235-264.

—— (2017b). Ryle on the Explanatory Role of Knowledge How. *Journal for the History of Analytical Philosophy* 5 (5).

—— (2019). Basic Action and Practical Knowledge. *Philosophers' Imprint* 19.

—— (2020). Practical Knowledge and Habits of Mind. *Journal of Philosophy of Education* 54 (2):377-397.

—— (2021). The Intelligence of Virtue and Skill. *Journal of Value Inquiry* 55 (2):229-249.

—— (2022). The Practicality of Practical Inference. In Adrian Haddock & Rachael Wiseman (eds.), *The Anscombean Mind*. New York, NY, USA: pp. 253–290.

Soteriou, M. (2009). Mental agency, conscious thinking, and phenomenal character. In Lucy O'Brien & Matthew Soteriou (eds.), *Mental actions*. New York: Oxford University Press. pp. 231.

Spaulding, S. (2016). Imagination Through Knowledge. In In A. Kind & P. Kung (Eds.), *Knowledge through Imagination* (pp. 207–226). Oxford: Oxford University Press.

Springle, A. (2019). Perception, Representation, Realism, and Function. *Philosophy of Science*, 86(5), 1202-1213.

—— (2021). Intentionality: A Problem-Solving Approach. Dissertation.

Springle, A. and A. Buccella (2024). Practical perceptual representations: a contemporary defense of an old idea. *Synthese* 203 (3):1-18.

Springle, A. and S. Goldwasser (under review). Episodic Memory as First-Hand Historytelling.

Springle, A. and J. Humphries (2021). Anti-intellectualism, instructive representations, and the intentional action argument. *Synthese* (3):7919-7955.

Stanley, J. and T. Williamson (2001). Knowing How. *Journal of Philosophy* 98 (8):411-444.

Steward, H. (2009). "Sub-Intentional Actions and the Over-Mentalization of Agency." *In New Essays on the Explanation of Action*, eds. Constantine Sandis, 295–312. Palgrave Macmillan

Strawson, G. (2003). Mental ballistics or the involuntariness of spontaneity. *Proceedings of the Aristotelian Society* 103 (30): 227-257.

Szpunar, K. K., J. M. Watson, and K. B. McDermott (2007). Neural substrates of envisioning the future. *Proceedings of the National Academy of Sciences* 104(2), 642–647.\

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, 120(3), 439–471. https://doi.org/10.1037/a0033138

Thompson, R.F. (1991). Are memory traces localized or distributed? *Neuropsychologia*. 29(6): 571-82.

Thompson, M. (2008). *Life and Action: Elementary Structures of Practice and Practical Thought.* Harvard University Press

Tolkien J. R. R. (1993). *The return of the king: being the third part of The lord of the rings. Second edition*. Boston, Houghton Mifflin Co.

Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson, *Organization of memory.* Academic Press. 381-403.

—— (1983). *Elements of Episodic Memory*. Oxford University Press.

—— (1985). Memory and consciousness. *Canadian Psychology / Psychologie canadienne* 26(1), 1–12.

—— (2002). Episodic Memory: From Mind to Brain. *Annual Review of Psychology* 53: 1–25.

Van Leeuwen, N. (2011). Imagination Is Where the Action Is. *Journal of Philosophy*, 108(2), 55–77.

—— (2013). The meanings of "imagine" part i: Constructive imagination. *Philosophy Compass* 8(3), 220–230.

—— (2016). The Imaginative Agent. In A. Kind & P. Kung (Eds.), *Knowledge through Imagination* (pp. 85–109). Oxford: Oxford University Press.

Velleman, D. (1989). *Practical Reflection*. Princeton University Press.

Vendler, Z. (1957). Verbs and times. *Philosophical Review* 66 (2):143-160.

Verwey, W. B. (2010). Diminished motor skill development in elderly: Indications for limited motor chunk use. *Acta Psychologica* (amst) 134 (2): 206–214. https://doi.org/10.1016/j.actpsy.2010.02. 001.

Verwey, W.B., E.L. Abrahamse, M.F.L. Ruitenberg, L. Jiménez, and E. Kleine (2011). Motor skill learning in the middle-aged: Limited development of motor chunks and explicit sequence knowledge. *Psychological Research Psychologische Forschung* 75 (5): 406–422. https:// doi.org/10.1007/s00426-011-0320-0.

Vierkant, T. (2018). Choice in a two systems world: picking & weighing or managing & metacognition. *Phenomenology and the Cognitive Sciences* 17 (1):1-13.

—— (2022). *The Tinkering Mind*. Oxford, UK: Oxford University Press.

Vitz, R. (2008). Doxastic voluntarism. *The Internet Encyclopedia of Philosophy*, ISSN 2161-0002, https://iep.utm.edu/

Volger, C. (2002). *Reasonably Vicious.* Harvard University Press.

Wagner, I. C. (2016). The integration of distributed memory traces. *The Journal of Neuroscience* 36(42), 10723–10725.

Wagoner, B. (2017). *The Constructive Mind: Bartlett's Psychology in Reconstruction*. Cambridge University Press.

Walton, K. (1978). Fearing fictions. *Journal of Philosophy* 75 (1): 5-27.

—— (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press

Weinberg, J. M. and A. Meskin (2006). Puzzling over the imagination: Philosophical problems, architectural solutions. In S. Nichols (Ed.), *The Architecture of the Imagination: New Essays on Pretence, Possibility, and Fiction*, pp. 175–202. Oxford University Press.

Werning, M. (2020). Predicting the Past from Minimal Traces: Episodic Memory and its Distinction from Imagination and Preservation. *Review of Philosophy and Psychology* 11(2): 301-333.

Williams, D. (2021). Imaginative constraints and generative models. *Australasian Journal of Philosophy* 99(1), 68–82

Williamson, T. (2007). *The Philosophy of Philosophy.* New York: Blackwell

——(2016) Knowing by Imagining. In A. Kind & P. Kung (Eds.), *Knowledge through Imagination* (pp. 113–23). Oxford: Oxford University Press

Wiltsher, N. (2019a). Characterizing the Imaginative Attitude. *Philosophical Papers* 48 (3): 437-469.

——(2019b). Imagination: A Lens, Not a Mirror. *Philosophers' Imprint* 19 (30): 1-30.

——(2023). Imagination as a process. Philosophy and Phenomenological Research, 106, 434–454. https://doi.org/10.1111/phpr.12861

Windt, J. M. (2010). The immersive spatiotemporal hallucination model of dreaming. *Phenomenology and the Cognitive Sciences* 9(2), 295–316.

Winlove, C. I. P., F. Milton, J. Ranson, J. Fulford, M. MacKisack, F. Macpherson, and A. Zeman (2018). The neural correlates of visual imagery: A coordinate-based meta-analysis. *Cortex* 105, 4–25.

Wittgenstein, L. (1980). *Remarks on the Philosophy of Psychology: Volume II*. Translated by C. G. Luckhardt and M. A. E. Aue. Edited by G. H. Von Wright and H. Nyman. Oxford: Blackwell.

Wolfson, B. (2016). Second Nature and Basic Action. In Roman Altshuler & Michael Sigrist (eds.), *Time and the Philosophy of Action.* New York, USA: Routledge. pp. 52-66.

Wu, W. (2011a). Attention as selection for action. In C. Mole, D. Smithies, and W. Wu (Eds.), *Attention: Philosophical and Psychological Essays*, pp. 97–116. Oxford University Press.

—— (2011b). Confronting many-many problems: Attention and agentive control. *Noûs* 45(1), 50–76.

—— (2013a). The case for zombie agency. *Mind* 122(485), 217–230.

—— (2013b). Mental action and the threat of automaticity. In A. Clark, J. Kiverstein, and T. Vierkant (Eds.), *Decomposing the Will*, pp. 244–61. Oxford University Press.

—— (2014). *Attention*. Routledge.

—— (2016). Experts and deviants: The story of agentive control. *Philosophy and Phenomenological Research* 92(2), 101–26.

—— (2019). Action always involves attention. *Analysis* 79(4), 693–703.

—— (2023a). Attending as a mental action. In Michael Brent & Lisa Miracchi (eds.), Mental Action and The Conscious Mind. Routledge. pp. 61-78.

—— (2023b). *Movements of the Mind: A Theory of Attention, Intention and Action*. Oxford: Oxford University Press.

Wymbs, N.F., D.S. Bassett, P.J. Mucha, M.A. Porter, and S.T. Grafton. (2012). Differential recruitment of the sensorimotor putamen and frontoparietal cortex during motor chunking in humans. *Neuron* 74: 936–946. https://doi.org/10.1016/j.neuron.2012.03.038.

Yablo, S. (1993). Is Conceivability a Guide to Possibility? *Philosophy and Phenomenological Research*, 53(1), 1–42.