

UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

Body and Memory: An Inquiry Into Skill

Permalink

<https://escholarship.org/uc/item/6c95615m>

Author

Herrmann, Rotem

Publication Date

2023

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Body and Memory: An Inquiry Into Skill

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Philosophy

by

Rotem Herrmann

September 2023

Dissertation Committee:

Dr. Eric Schwitzgebel, Chairperson

Dr. Luca Ferrero

Dr. Howard Wettstein

Copyright by
Rotem Herrmann
2023

The Dissertation of Rotem Herrmann is approved:

Committee Chairperson

University of California, Riverside

ACKNOWLEDGEMENTS

I wish to thank, first and foremost, my committee, Eric Schwitzgebel, Luca Ferrero, and Howie Wettstein, for all their help and flexibility, especially in the last stages of writing the dissertation. I am greatly indebted to them for allowing me to work in the rather silly and unreasonable manner I did and working with me in a way that goes above and beyond the call of duty.

I would also like to thank my dissertation writing work group: Elliot Koch, Shamoni Sarkar and Josh Waugh. Our weekly meetings kept me sane and were often highlights of the writing process. Their help and friendship got me through the early stages of writing in ways I would never have been able to on my own.

Thanks also goes to my partner, Chris, my parents, and many dear friends who supported me from the start. My beloved younger sibling, Ronnie, deserves particular mention for all their unflinching support over the years and for being a bright light in my life during the darkest parts of this journey (and long before that, too).

My cats, One, Two and Momo, were beacons of love and cuddles. They kept me company during every stage of writing, and kept my days beautifully fluffy, funny and weird.

Lastly, I want to thank the Pomodoro Study Playlist on the ASMRWeekly YouTube channel for being the soundtrack and engine behind any work I've managed to do in the last few years since I found it. This project owes the channel's creator a great debt.

I dedicate this dissertation to myself.
I wouldn't have done it for anyone else, and I couldn't have done it without me.

ABSTRACT OF THE DISSERTATION

Body and Memory: An Inquiry Into Skill

by

Rotem Herrmann

Doctor of Philosophy, Graduate Program in Philosophy
University of California, Riverside, September 2023
Dr. Eric Schwitzgebel, Chairperson

This project aims to produce an account of procedural memory, specifically as it relates to physical skills. Part 1 is an exploration of the nature of skill moving through literature spanning philosophy of mind, psychology, cognitive science and epistemology. From this exploration, several key criteria for skill are developed that will need to be met and maintained in an account of the memory of this phenomenon throughout the rest of the project. Building on these criteria, Part 2 demonstrates that existing dominant models of memory, namely Storage and Retrieval, Reconstructive, and Simulationist accounts, are not sufficient to encapsulate or explain procedural memory, though they do offer some hints as to what we should be looking for. Based on the shortcomings of existing accounts, Part 3 examines accounts from an altogether different approach, namely embodied views on memory. However, it is determined that not all embodied views are fit for an account of procedural memory and that views that lay out a clearer connection with the experience of time and bodily influence in memory are necessary. As such, the project ultimately puts forward a positive account inspired by figures like Merleau-Ponty and Maxine Sheets-Johnstone whereby procedural memory is a process of holding together one's past, present and future through intentional action and familiar kinesthetic movements.

TABLE OF CONTENTS

Part 1: A Walk Through of Motor Skill	1
Chapter 1: Accounts of Skill in Philosophy of Mind, Cognitive Science, and Psychology.	4
Chapter 2: Epistemological Accounts of Skill	57
Chapter 3: Walking as Skill	104
Part 2: Memory and Skill	141
Chapter 4: Storage and Retrieval Model of Memory	155
Chapter 5: (Re)Construction and Simulation Models of Memory	183
Part 3: Body and Memory	221
Chapter 6: Embodied Views of Memory	224
Chapter 7: Body and Time in Memory	257
References	295

Part 1:

A Walk-Through of Motor Skill

Introduction:

When we think of what it means to remember something, we usually think of things like remembering your partner's birthday, or the day you graduated from high school, or that Paris is the capital of France, or that Robert the Bruce led the Scots to independence against the English in 1314. That is, we usually think of memory as having to do with remembering facts, figures, events from history or one's own life. In this sense, what we usually think of as memory has to do with remembering what we colloquially think of as 'information'. But there are other things I might be said to remember. For instance, suppose (completely hypothetically, of course) I learned how to play clarinet as a child, played it regularly till the end of high school and then never touched it again. If I was to pick up a clarinet today and manage to play it competently, we might be inclined to say something like, "hey, whaddya know! She remembers how to play!". Whatever I might be said to remember, in this case, is what psychologists call procedural memory, or what philosophers tend to call non-declarative memory. And whatever I'm remembering, seems to be importantly different than the 'information-based' stuff that we mentioned earlier.

So, what is this procedural memory, then? The category of procedural memory is usually thought to cover several phenomena, including priming effects and habit, but one of its primary domains is understood to be that of skill¹: this is the aspect of procedural

¹ More on this in Part 2, Chp 4.

memory I am most interested in. The aim of this project is to get a grip on how memory plays a role in physical skill and skilled behavior, like that of playing an instrument or a sport. To do this, we will turn our attention to leading models of memory and how it works, and determine how well they apply to our procedural case. The first of these will be Storage and Retrieval (S&R) views (Chp 4), followed by more contemporary Reconstructive views (Chp 5). However, we will see that neither of these accounts will be able to help us, largely because they are built around the former ‘information-based’ type of memory. Given the failure of these previous views to account for our target phenomenon, and since our focus is on physical skills for which the body is presumably integral, we will then turn our attention to some alternative understandings of memory: namely, Embodied views (Chp 6) and Phenomenological views (Chp 7). Using the phenomenological views, especially that of Merleau-Ponty, we will end our investigation with an account of procedural memory that is fundamentally different from many of our original starting points and from our usual understandings of what it means to remember something. Ultimately, we will see that, in order to explain procedural memory, we will need to tie the remembering subject in with their environment, with their positions in space and time, and with their bodies as the intersection of all these dimensions. All of this will mean that memory is not so much a question of bringing past information or experiences to bear on your behavior, but of collecting and holding your entire subjectivity together across time and action. But we are still far from that...

However, since skill is so central to this kind of memory, it will be crucial to get a clear picture on what ‘skill’ is and what it means to be skilled at something. With this

task in mind, Part 1 of this dissertation consists of an exploration of skill – how it is characterized across various disciplines, its acquisition, its relation to knowledge, its relation to other phenomenon like habits, etc. Chapters 1 and 2 begin this exploration and involve extrapolating key criteria for skill from literature in Philosophy of Mind, Cognitive Science and Psychology on the one hand (Chp 1), and Epistemology on the other (Chp 2). In Chapter 3, I address a key example of the sort of skill I see as central for much of what happens in later chapters of the dissertation, namely walking. Walking is understood by many to be a ‘basic’ skill – one exemplary of the foundations of bodily engagement in skillful action and one that lays the grounds for many other skills. But many others understand walking as ‘too basic’ to be a skill – they see it as an innate ability and so out of the running for consideration as a skill. I demonstrate that, contrary to some intuitions, walking is in fact a skill and that it meets all the standard criteria and understandings of skill developed in the earlier chapters.

Chapter 1:

Accounts of Skill in Philosophy of Mind, Cognitive Science, and Psychology

1.0: Introduction:

In trying to get a grip on what it means to have a skill, we will first begin with an investigation about the nature of skilled performance. For instance, suppose you know how to play volleyball, and suppose you're quite good at it. Suppose further that someone who wasn't as skilled as you wanted to know about how you perform so well. They might ask things like, "How did you become so skilled? Was it something you were innately and immediately good at? Or was it something you learned over time? When or how did you realize you were getting better? What's it like when you play? Are you thinking about everything you're doing? Or are you in a 'flow-like' state, just letting things happen? How do you keep track of everything going on?"

These are precisely the kinds of questions asked in more empirically minded investigations across philosophy of mind, psychology and cognitive science, all of which we will explore in this chapter. We will see that there are generally two main camps, and so two main sets of tendencies, in answering these questions: we will call them 'intellectualists' and 'non-intellectualists'. Intellectualists tend to give answers that highlight the cognitive control and effort of the performer in being able to perform as well as they do. Non-intellectualists, on the other hand, tend to emphasize the automaticity and 'flow' of their actions as foundational for their performance. However, will see that, though both sides have evidence in their favor, neither is able to fully take account of the other's. As such, taking on a hybrid position between these two camps

seems the best way to account for as much evidence as possible, and should serve as the model for the rest of our investigation.

1.1: Learned vs Innate:

In order to characterize skill, most accounts in philosophy, psychology and cognitive science focus on several key aspects of the phenomenon. These aspects are intuitively, phenomenologically and empirically backed, in most cases. One of the key features of skill is that it is learned – in other words, it cannot be an exercise or action that is innate to the individual². For example, one wouldn't usually expect that an activity such as breathing or digesting (at least under normal circumstances) amounted to any kind of skill – it's simply an innate capacity of a human body. On the other hand, the classic examples of skill (such as riding a bike, playing a sport or an instrument, etc.) are activities that an individual has to learn over time. Though some people may be more 'innately' or 'naturally' disposed towards these various activities, it is thought that performing with any degree of competence, or successful performance more generally, will always require some amount of learning and practice.

Indeed, most accounts of skill focus on its acquisition and understand this to be at the core of the phenomenon. There are many such accounts that put the learning of a skill at its core. For example, (Stanley & Krakauer, 2013) say that “skill can be considered the practice-related improvement in a goal-oriented action” (pg 5). Similarly, Fridland (2014)

² Some have argued that it is not clear whether learning is a necessary feature of skill, but they agree that “it is reasonably clear that most human skill involves substantial learning”. (Christensen et al., 2019)

explains that skill³ is practiced, where practice is understood “to require, at least temporarily, attending to and attempting to improve an ability as an end in itself” (2732).

One of the most classic views of skill, namely Fitts and Posner’s account (1967), focuses primarily on skill acquisition and describes three stages of this process⁴. The first is what they call the *cognitive* stage, in which the practitioner performs the actions involved in the task but at a level that only approximates ‘proper’ performance. The second stage is the *associative* stage, in which the practitioner notices and corrects errors made in performance. The third and final stage is the *autonomous* stage, in which the practitioner starts making continuous, incremental improvement in their performance in such a way that demonstrates increasing ‘autonomy’ from cognitive control (Fitts & Posner, 1967).

As this is very widely accepted, our first criterion of skill will be as follows:

Criterion 1: Skill, in so far as it relates to the ability to successfully perform an activity, is learned, acquired or gained through practice. It is not innate ability.

1.1.1: Dreyfus and Dreyfus on Skill Acquisition

Even though the general idea described in Criterion 1 is very well accepted, the precise means or mechanisms of this learning is not as generally accepted. I want to focus on perhaps the most well-known and most referenced philosophical account of ‘skill as learned’ and its process of acquisition. This provides us with an example to work through

³ Fridland actually phrases this sentiment as ‘control’ being learned with practice. However, as will be seen below, Fridland understands control to be the core of skill and the aspect of it that we should be most concerned about. As such, I felt the shift in terminology in this instance would be acceptable when speaking in broad or general terms.

⁴ This approach has been labeled an ‘information processing theory’ (see (Toner et al., 2015) in which controlled processing eventually makes way for automatic processing.

more carefully and a reference point from which to make comparisons and critiques. The account in question is that put forward by (H. Dreyfus & Dreyfus, 1986), (henceforth D&D) in which they explain that skill is “acquired from practice and sometimes painful experience”⁵ (ibid, 16). D&D explain that, in some instances, it is acquired through trial and error or guided imitation of more seasoned practitioners. However, they suggest that it is more common for adults to begin “to acquire new skills by means of either written or verbal instruction” (ibid, 19) – this latter means of acquisition is the one they’re most interested in and use as a basis for exploring the means of skill acquisition in general.

D&D’s account posits five stages of learning acquisition and has proven influential in the literature at large. These five stages progress sequentially from one to the next and, according to D&D, have the advantage of explaining, (1) the transition from novicehood to expert over time and repeated practice and, (2) how the most ‘talented’ practitioners at later levels perform better than even those most talented practitioners at earlier levels (ibid, 21).

The first stage of D&D’s skill acquisition model is the ‘Novice’. Here, the practitioner learns various basic facts and features about the performance of the skill as well as some rules about how to incorporate those facts and features into their performance. D&D stipulate that there will be some elements of a situation that are so clearly defined for the novice that they can recognize them without reference to the overall situation – these they call ‘context-free’ elements. Further, they understand the manipulation of these ‘context-free’ elements based on rules to be ‘information

⁵ They equate skill with ‘knowledge-how’ – more on this in Chp 2.

processing'. 'Information processing' is understood to be distinct from more holistic methods of recognition, for example 'template matching'. In other words, the manner in which a novice proceeds is like "a computer following a program" (Dreyfus, 1997, 19) and involves a *detached* rule-following stance.

To play out these stipulations, consider the following example (which D&D use but do not flesh out fully). Consider a novice driver who is presented with information such as the speed at which to drive in a residential area; or what distance to keep between themselves and the car in front of them when pulling up behind it; or looking '10 seconds ahead' on the road while driving; etc. All of these facts are context-free, as they are universal or general rules rather than specific applications to a particular instance of driving and the rich context in which that instance would exist. For example, it doesn't consider the traffic conditions of the neighborhood, condition of the car, unanticipated stops, etc. Suppose a novice driver recognizes that the distance to maintain between themselves and a car in front of them when pulling up behind it is, say, a car's-length. If they do this based on the application of rules around this context-free element, for example, thinking to themselves 'if I am pulling up behind a stopped car on the road, then I should keep a car's distance between me and the car in front of me', then they would be doing so by 'information processing'. If, however, they recognize this based on some other measure, for example, that this is the distance they usually maintain while driving

and that it is the distance they have been taught in the past, they will be doing so based on a ‘holistic⁶ template matching’.

However, since they have not yet developed a more complete sense of the task at large, the novice judges their performance based on how well they follow the rules. That is, though they have a small pool of past experience to guide them, their judgment of their own performance will still be largely based on holding themselves to the handful of initial rules they learned. In order to improve, D&D suggest that the novice must acquire several such rules and accumulate experience around them to help them move on to the next stage. They also add that, at this stage, performance of the skill in line with all the rules they have acquired will require a great deal of concentration.

The second stage of D&D’s skill acquisition model is the Advanced Beginner stage. Here, after considerable experience in real-world situations, the practitioner advances their skill to a ‘marginally acceptable level’. While they do learn more context-free facts and use more sophisticated rules, “it also teaches [them] a more important lesson involving an enlarged conception of the world of the skill” (D&D, 1986, 22). That is, with more practical experience comes an increased ability to recognize elements of the skill on their own. D&D call these ‘situational’ elements, because they are based on

⁶ A point of clarification - one might worry that such a driver might still be determining where to stop in a more ‘information-based’ sense, since they may still be deducing from past experience where the appropriate place is to stop. However, the sense in which it would be more ‘holistic’ for D&D is that it relies on their grip of the situation at large rather than any kind of deduction or ‘analytic’ reliance on past experience – for instance, the nervousness they feel if they get too close to the car ahead, the ‘familiarity’ of starting to break at the appropriate distance given their understanding of the car, etc. In this sense, the driver is relying on a more holistic relationship between past experience and present circumstances to help guide them, rather than actively deducing the appropriate place to stop. But we should note that, at this stage, this kind of template matching will not be robust and will be based on relatively little past experience. Whatever the ‘holistic’ understanding is at this point should be understood to be only the first steps in the process described below.

particular, real-world experiences in a variety of contexts. D&D also suggest that this shift has to do with ‘perceived similarity with prior examples’. That is, at this stage, practitioners are just beginning to build up their own understandings and familiarity with the skill based on past experience and recognizing its relevance in present circumstances. With this level of experience, rules can now be applied to both situational and context-free components of the skill.

Consider the same example of learning to drive. The advanced beginner driver can now, according to D&D, use situational elements to help in their decision making about driving. For example, they might now be able to use engine sounds as an additional factor in their application of rules: if the engine is making an unusual sound once it reaches a given speed, the driver might decide not to accelerate any further or to slow down. The engine’s sound is a situational factor, one having to do with this particular instance of driving, that can now be incorporated into their decision-making process.

The third stage is Competence. At this stage, the practitioner can recognize a very large degree of context-free and situational elements in real-world circumstances. However, in the lead up to this stage, they do not yet have a sense of which of these are most important and so this recognition of elements can become overwhelming. Because of this, a mark of having achieved competence is to adopt a ‘hierarchical procedure of decision making’. These usually consist of, first, choosing a plan to organize the situation, and then examining all the elements and ranking them based on the chosen plan. With this in hand, the practitioner can then form ‘constellations’ of elements and learn that, when a given constellation arises, they should act in some particular way.

Unlike the novice, who was provided with a set of rules to follow and dutifully does so, the competent performer begins to organize and prioritize those rules using their own judgement and experience.

Back to the driving example, a competent driver is now no longer focusing solely on facts about how to drive generally, but will also be motivated by a situational goal, such as getting from A to B *quickly*, perhaps due to tardiness. With this in mind, the driver can select the most direct route with the least traffic, putting aside other features such as the familiarity of one route or the scenic beauty of another. Similarly, because they are driving in accordance with their goal, they might not focus so much on maintaining the same distances between themselves and the car in front, or on staying within the speed limit, etc; speed may require pulling up alongside a car so they can turn right rather than waiting behind it, for example, or simply driving faster than the speed limits of that area allow. This being the case, some of the elements that they would initially have intently focused on in previous stages can be deprioritized and put aside in favor of achieving larger goals.

Choosing a plan, however, is not as simple as merely following rules. There is no objective procedure on how to choose a plan and yet this choice is absolutely critical in the process of skillfully performing an action. In fact, D&D stipulate that performance at this level *requires* choosing a plan to organize one's behavior. They explain that the "combination of nonobjectivity and necessity introduces an important new type of relationship between the performer and [their] environment" (D&D, 1986, 26). Specifically, the choice of a plan introduces the notion of *responsibility* for said plan and

its consequence, which brings with it emotional engagement with one's choices. So, while the choice of plan is made in a detached manner, the practitioner will ultimately find themselves deeply *involved* in the results of their actions. A successful outcome will bring deep satisfaction, while failures will be painful. This involvement is actually crucial for future development because "resistance to the frightening acceptance of risk and responsibility can lead to stagnation and ultimately to boredom and regression" (Dreyfus, 1997, 21).

The fourth level of skill acquisition is Proficiency. One of the new features of this level of skill is a kind of deep involvement in the task that establishes a responsive perspective for the practitioner. Once in this deeply involved perspective, certain parts of a situation will stand out to the practitioner while others will be ignored. Experience will be largely responsible for what stands out and what doesn't – whatever the proficient practitioner has learned is most relevant will pop out, while anything less than relevant will remain 'in the background'. This highlighting of salient features happens naturally and intuitively for the practitioner – they are not calculating or determining that they are relevant, they already stand out as such given their past experience.

Throughout the course of action, these salient features, plans, expectations, etc. will change as the conditions of the activity change and the proficient practitioner will have to make adjustments to their behavior. Choices are made in this state, but they are made without conscious deliberation. Rather than applying rules to salient elements to build up a holistic understanding of the situation, the proficient practitioner is already engaged with the holistic pattern and reacts accordingly. D&D suggest that this is

because “the proficient performer has experienced similar situations in the past and memories of them trigger plans similar to those that worked in the past and anticipations of events similar to those that occurred” (D&D, 1986, 28)⁷. They call this ‘intuitive’⁸ ability ‘holistic similarity recognition’ (ibid). In other words, as the practitioner becomes increasingly ‘involved’ in performance, they will start to develop holistic understandings of the patterns of a task and “intuitive behavior gradually replaces reasoned responses” (Dreyfus, 1997, 21). Once this intuitive understanding is established, they will be able to “simply see what needs to be achieved rather than deciding, by a calculative procedure, which of several possible alternatives should be selected” (ibid). As this happens, action becomes easier and less stressful, and one experiences less doubt about one’s actions – indeed, “at the moment of involved intuitive response there can be no doubt, because doubt comes only with detached evaluation of performance” (ibid).

Though some parts of the proficient performer’s action will operate on this intuitive level, including the ways in which they organize and understand their task, D&D explain that they will still find themselves thinking analytically about other aspects of the task. Namely, certain parts of their planning, including how best to manipulate the environment. This is because the proficient practitioner has enough experience to guide

⁷ D&D also note that there is an element of comfort or discomfort that comes with the recognition of past experiences. If an activity is lining up neatly with past experience, the practitioner will likely feel a degree of comfort in making their decisions. Conversely, past experiences do not line up neatly with the current circumstances, this will likely cause the practitioner to feel discomfort, which will drive them to search for options that reduce their unease while maximizing desirable outcomes. (D&D 1986, 37).

⁸ D&D understand ‘intuition’ to involve ‘deep situational involvement and recognition of similarity’ (29) and treat it as synonymous with knowledge-how (28). It is nonconscious and non-inferential, but it is not simply guessing, divine inspiration, or the like. They believe the basis of intuition is the memory of whole past situations and what happened after them, and this provides the basis of *intuitive expectations* (38).

their intuitive grip on goals and important features of a situation (i.e. a big picture), but not yet enough to intuitively select the best response to every scenario or how best to utilize environmental factors (i.e. finer points of the big picture). As such, the proficient practitioner may still *decide* what to do by falling back on detached, rule-based processes. For example, a driver approaching a tight curve might intuitively realize that they are driving too fast to go through it safely. Once the ‘spell’ of intuition is broken, they will then make an analytic decision about whether to take their foot off the accelerator, apply the brakes, etc. Thus, analytic decision making still happens at this level, but it is separate from the truly ‘intuitive’ aspects of performance of the skill.

Finally, the fifth stage of skill acquisition is Expertise. Like the proficient practitioner, the expert is deeply involved in their performance of the skill and intuitively knows *what* needs to be achieved based on a deep and practiced understanding. However, unlike the proficient practitioner, the expert also knows *how* to achieve their goals on an intuitive level – this is the distinguishing mark of an expert. This is achieved by the expert developing a more refined sense of discrimination than the proficient practitioner: looking at the big picture, the expert is able to distinguish between situations requiring one response from those requiring another. This is gradually acquired with experience and the application of the same involvement and intuitive response down to these ‘lower’ levels (Dreyfus, 1997, 22). Much of their ‘intuitive’ performance will have to do with memory ‘triggering’ similar past experiences and acting in accordance with what worked for them in the past. Since the expert does not need to ‘detach’ to consider plans or how to execute them, they undergo complete immersion in a task and do not separate

themselves from it. That is, our driver would ‘become one’ with the car – rather than understanding themselves as ‘driving a car’, they are simply ‘driving’ (D&D, 1986).

Usually, the expert will not make decisions based on analytical calculations or formulaic rules – they act based on their experience of their position in the activity and how it relates to their past experience (ibid, 37). Their experience, itself, indicates when something feels amiss and a new course should be taken, and course-corrections will happen in the same ‘intuitive’ way described above. D&D explain that “[w]hen things are proceeding normally, experts don’t solve problems and don’t make decisions; they do what normally works” (original emphasis, ibid, 30-31). However, just like in the previous level, there may be instances in which the expert will have to ‘think’. In situations where the outcomes are particularly important and when the expert has enough time, they may deliberate about how to act. However, “this deliberation does not require calculative problem solving, but rather involves critically reflecting on one’s intuitions” and will happen outside of the ‘flow’ state (ibid, 32).

The picture D&D paint, then, is one where skill is gradually acquired and moves from rigid rule-following to expert intuitive ‘flow’ – from detached top-down control to involved equilibrium. As previously mentioned, this account has proven extremely influential and serves as one of the central philosophical accounts of skill acquisition. However, one of the reasons it is considered so central is because it is perhaps the quintessential ‘non-intellectualist’ understanding of skill and serves as the main target for many intellectualist positions. As this and the next chapter proceed, I will explore Hubert Dreyfus’s further developments on this view and give further insight into mechanics of

what he often refers to as ‘skillful coping’, i.e. fluid, ‘unminded’ engagement in an environment. Similarly, I will explore the many points of disagreement that intellectualist positions take up with his position.

First, however, there seems to me to be a common set of concepts that will be useful to lay out, as most views agree that these features are characteristic of skill and try to account for them in different ways, namely fluency and flexibility. With these in hand, we will be better equipped to flesh out both intellectualist and non-intellectualist understandings of skill.

1.2. – Fluency and Flexibility

In addition to being learned, two core features of skill described across the literature are fluency and flexibility⁹. ‘Fluency’ of skill has to do with the smooth and graceful manner in which skilled action is performed. Novices, for example, will often perform tasks in rather faltering or disconnected ways that betray their status as novices. Skilled practitioners, on the other hand, will often perform tasks with fluidity and grace. For example, Adolph & Robinson, (2013) describe the fluency of skill as follows: “Fluency is what makes skills efficient, coordinated, and beautiful to observe. It is the ability to execute movements smoothly, accurately, and rapidly” (14). Building on ideas from D&D, fluency (or in their parlance, fluidity), involves complete immersion in a task such that the practitioner seemingly operates completely intuitively and seamlessly.

⁹ The terminology used to describe these same concepts varies across the literature, but I think this wording is accessible and intuitive.

Flexibility, on the other hand, has to do with the practitioner's ability to adapt the performance of their skill across varying circumstances – that is, their ability to adequately respond to the various changes, pressures and difficulties of a situation and still perform well. Part of the idea here is that any given skill will have to be applied over an extraordinarily large array of circumstances – no two instances of performance are ever exactly the same. As such, a skill must exhibit flexibility by being deployable and consistent across all these instances. The flexibility of a skill might entail being applied across all the various instances within set or 'normal' parameters of the skill (for example, being able to play basketball on indoor courts under the 'normal' conditions of time, other players, the usual equipment, etc.), or it might entail being applied across increasingly dissimilar parameters (for example, playing outdoor basketball where temperature and wind might have to be taken into account, or playing with a much heavier/lighter ball, etc.). Being able to use skill across these instances would demonstrate the flexibility of a skill.

These parameters are also very widely accepted and so our second criterion of skill will tentatively be...

Criterion 2: skill is typically manifested in fluent and flexible action.

This much is uncontroversial. However, there is extensive literature on the nature of fluency and flexibility and what kind of processes help instantiate them. More specifically, there is a debate between what we might call intellectualists and non-intellectualists about whether fluency and flexibility are instantiated through cognitive control or through automatic processes, respectively. As such, in order to clarify our

second criterion and in order to reach our third, we will explore this debate more carefully. Through this investigation, we will see not only how each camp accounts for the fluency and flexibility of skill, but also the (often empirical) evidence they rally in their favor. Rather than ultimately offering a full endorsement of either of these positions, we will see that neither is able to fully capture or account for the strengths of the other. Instead, I will present an example of a commendable third-option view that maintains the strengths of ‘both sides’ and will suffice for our current purposes of understanding fluency and flexibility until we reach later stages of the dissertation.

1.2.1 – Non-Intellectualist Accounts of Skill

Non-intellectualist accounts aim to describe skill in a way such that it does not rely on any ‘thinking’ processes – these accounts vary across the literature and will often focus on different iterations of what kind of ‘thinking’ process they are interested in. Some aim to deny the role of explicit, conscious thought in performance, while others deny the role of any propositional thinking during performance, and others still deny the necessity of ‘representations’ for performance (D&D 1986 being a paradigm example, and Dreyfus’ later work covers all of the above). Much of this denial gets showcased in the non-intellectualists’ push for ‘automaticity’ of skill. For example, many people would acknowledge on both intuitive and phenomenological levels that when one first begins learning how to do something new, its performance is difficult and requires a great deal of attention. However, as one gets better at this activity, it’s commonly thought that performance becomes increasingly ‘automatic’ – i.e. the task or activity becomes easier and one can perform it without having to think so consciously about what one is doing. In

this context, automaticity involves what some have called ‘mindlessness’ or a ‘flow’ state – these are thought to involve both a lack of thought (conscious or otherwise) about how one is performing the task as well as a certain ‘fluidity’ of engagement.

Non-intellectualists will often cite automaticity as a means of encapsulating and explaining the characteristic features described above, i.e., fluency and flexibility. For example, fluency is often tied to the idea of automaticity in so far as it seems to come part and parcel with automaticity – that is, the more automatic one’s performance, the more fluent. Similarly, D&D (1986) describe the fluidity¹⁰ of performance that comes with expertise. As we saw, for them it has to do with the deep immersion in a task which means that the expert does not have to consciously think about how to act. For D&D, fluidity in skill comes when the practitioner is so expert that they can recognize what to do and how to do it based on a holistic understanding of their situation and do not have to pause, stop, or start to deliberate about this. They simply ‘react’ because they are so deeply bound up in the situation and recognize patterns that they have already encountered. As such, they can act in the quick and continuous ways that typify fluidity of action. Again, this fluidity is thought to be related to the automaticity of skill because, for actions to unfold so skillfully and so quickly, it is argued that there simply cannot be enough time for conscious deliberation to occur.

Non-intellectualist accounts also understand automaticity to be involved in the flexibility of skill. There are at least two ways that flexibility is thought to be instantiated in automated actions (Christensen et al, 2019): first, the performance of a skill might

¹⁰ I’m understanding ‘fluency’ and ‘fluidity’ to be synonymous and interchangeable.

require the flexible use of general structures or patterns of behavior. For example, a soccer player would have general structures or patterns of behavior in place that relate to kicking the ball which they must flexibly perform depending on the particular goings on of the game at that moment. At one point in the game, they may have to kick lightly enough to dribble and, at others, they may have to kick and swing precisely so as to make a corner shot, etc. Here, then, flexibility has to do with modifying the particulars of a general automated pattern of behavior.

The other way flexibility is instantiated in automatic behavior is in the flexible organization of actions in performing a task. That is, there may be several actions that a practitioner must undertake to achieve a particular goal, but those actions can be performed in a flexible order that still culminates in the achievement of the same goal. For example, when cooking a meal, though there are some steps that must happen in a particular order (such as when particular foods get added into the skillet, etc.), there are other facets of the performance that can be performed in any number of orders. The cook could prepare a *mise en place* and gather all their ingredients before starting to cook, or they might chop and prepare ingredients in their order of addition as earlier ingredients are already the process of cooking. Furthermore, within either of those choices, the order in which they choose to prepare multiple ingredients can happen in any number of orders (though it may be more limited in the latter case). Here, then, flexibility is instantiated in the order in which ‘smaller’ actions are carried out within a larger automated process.

Since non-intellectualists argue that automaticity is linked to higher levels of fluency and flexibility, they also suggest it is linked to increasing levels of skill. That is,

non-intellectualists often describe the earlier stages of skill acquisition as cognition-heavy while later, more expert stages, are increasingly ‘automatic’. For instance, recall (Fitts & Posner, 1967), described briefly in an earlier section. Their view describes a three-stage transition from cognitive to associative to autonomous learning/use of skill. On this account, skill becomes increasingly ‘autonomous’ from cognitive control as it becomes increasingly competent. Similarly, Dreyfus and Dreyfus (1986), who are among the strongest proponents of the automaticity of skill, offer an account in which thought-heavy problem-solving may be sufficient for producing skilled behaviors but is not necessary. On their view, the acquisition of skill involves a transition from thought-heavy problem-solving to intuitive, fluid responses based on a holistic responsiveness to one’s context. The more skilled one gets, the more ‘automatic’ its deployment due to being more deeply ‘in-tune’ with one’s context.

Recall the transition up the five stages of skill acquisition described above – D&D suggest that lower levels (novice, advanced beginner, and competence) will involve some level of conscious, problem-solving decision making. However, the latter stages (proficiency and expertise) are “characterized by a rapid, fluid, involved kind of behavior that bears no apparent similarity to the slow, detached reasoning of the problem-solving process” (27). In other words, while acting skillfully, they deny that experts or highly skilled people make decisions about how to act based on creating and analyzing a list of potential actions¹¹, nor do they need to generate and execute long lists of actions that add

¹¹ One might worry that these kinds of ‘decisions’ or ‘calculations’ could be happening unconsciously and so bypass some of the non-intellectualist’s concerns. However, as some of the following points hint towards, even the unconscious versions of those thoughts would likely depend on some manner

up to a procedure that they execute point by point, nor do they make rules or heuristics and simply follow them. D&D are adamant that the expert doesn't follow any rules or make any calculations because no rules or calculations could encapsulate their skill. Instead, they learn to recognize and respond to the patterns of a given context intelligently but automatically. Indeed, D&D suggest that the slow, conscious, and intentional kind of deliberation about how to act found in lower levels is not representative of the way we usually go about in the world – these are the exceptional cases and only happen seldomly.

One example they use to help substantiate this understanding of the separation between cognitive or analytic ability from those involved in skill is that of International Chess Master, Julio Kaplan. D&D performed an experiment with him where he was asked to add heard numbers at a rate of $\sim 1/\text{second}$ while at the same time playing a game of five-second-a-move chess with another master player. D&D report that, even though Kaplan's analytic capacities were fully loaded during the performance of these tasks, he was still able to play very coordinated and skillful games of chess. D&D interpret this as evidence that skill cannot be primarily cognitive and must rather be based in automatic execution of tasks based on extensive past experience. (D&D 1986, 33).

Beyond accounting for fluency and flexibility, automaticity is also thought to capture some central experiences across skill performance that are repeatedly reported in

of propositions, lists, representations, etc. which they non-intellectualists deny could be in play, at all, even on the unconscious level. However, for a more detailed explanation for how a non-intellectualist, particularly Dreyfus, aims to account for this kind of automaticity without representations, etc., see Chp 2.

skill literature. As such, non-intellectualists understand these oft-reported experiences as support for their view. Christensen et al. (2016) list these experiences as follows:

1. Skilled performance involves *reduced attention*, or the decreased need to pay attention to performance once a skill has been acquired. That is, a skilled practitioner does not need to pay attention to the mechanics of performance in the same way a novice does, and so the amount of attention required to perform reduces once a skill is acquired.
2. *Multi-task tolerance* in which a skill that has been well-learned can often be performed simultaneously with another activity without sacrificing much by way of performance. For example, when one is learning to drive, it might be difficult to drive at the same time as conversing and also change the radio station. However, once one becomes skilled at driving, one can often drive at a high level of performance while simultaneously conversing and/or channel surfing the radio.
3. *Disruptive attention* in which paying attention to the performance of a skill can be detrimental to said performance. This is also sometimes referred to as ‘choking’, or ‘the yips’, wherein a practitioner is thought to be directing their attention away from a ‘flow’ and back to more ‘basic’ mechanics of performance. This is thought to have the effect of either making the practitioner unable to perform at all or unable to perform at their usual levels.
4. *Reduced cognitive effort* in which the practitioner reports a low sense of cognitive effort in performing a task. For example, learning a skill often

requires one's full cognitive capacity while practicing and is, as such, cognitively effortful. But as one acquires the skill, cognitive effort can be expended much less taxingly – performance no longer requires so many cognitive resources.

5. Finally, *reduced memory*, in which the practitioner often has low or reduced memory of having performed the skill. A common example of this is that, once one has become skilled at driving and has, for example, become accustomed to driving home one particular way, it is not uncommon for the driver to forget (parts of) the drive home – they might arrive in their driveway and be surprised to find themselves there. The reasoning here is that their performance was so fluid and engaged that their skill was deployed on 'auto-pilot' and required little guidance or monitoring to achieve the task. As such, little (if any) of the trip is recorded in their memory.

Furthermore, there are some empirical studies that help substantiate the non-intellectualists' support of automaticity in skill. For example, when practitioners first start learning a new motor skill, their motor cortex is in heavy use throughout performance. However, as they get better, activity in their motor cortex decreases and is gradually transferred to subcortical regions, especially the basal ganglia (Aldridge & Berridge, 2003; Poldrack et al., 2005). Similarly, as practitioners become better versed at motor skills, they begin to be able to sort their actions into 'chunks', meaning they can group a series of component actions together into one cohesive action. Being able to 'chunk' in

this way means they can perform the same task with decreasing demands on their attention (Boyd et al., 2009).

1.2.2 – Flow states - Thinking as detrimental to automatic processes:

Point (3) above, disruptive attention, suggests that paying (too much) attention to the finer points of a skill one is performing, e.g., starting to think about it, will be detrimental to said performance. This phenomenon has been identified as ‘choking’ or ‘the yips’ and the common wisdom suggesting ‘don’t think, just do’ can be found in all manner of writing on skill, including D&D, Fitts and Posner, the Zhuangzi (Ivanhoe & Norden, 2005) and writings on baseball by evolutionary biologist Stephen Jay Gould (Gould, 2004).

An oft discussed example of this kind of hinderance is the story of New York Yankee basemen Chuck Knoblauch. Knoblauch developed a series of throwing problems that seemed to indicate a lack of control over his performance – sometimes he was barely able to toss the ball and other times threw it shockingly far out of bounds. His impairment was popularly understood to be one of having started ‘thinking’ about his performance and suddenly starting to pay attention to what had been up till then his ‘automatic’ throwing processes. Having done this, he had significantly impaired his performance and was seemingly unable to return to the more ‘natural flow’ of his play. This case is often seen as emblematic of failing to adhere to the general wisdom, “don’t think, just do” – by getting too caught up in ‘thinking’, Knoblauch couldn’t ‘do’. As such, non-intellectualists suggest that, when performing well, practitioners will be performing in what is often called a ‘flow state’ which avoids active reflection on one’s performance.

A now familiar view helps account for this kind of conventional wisdom around flow states. Within their larger framework, D&D also describe a particular area of interest within ‘automatic’ performance which they call ‘flow’. D&D understand ‘flow’ states to be those in which all monitoring of behavior is suspended, performance reaches its peak, and which are accompanied by a feeling of euphoria. D&D explain that, rather than understanding ‘flow’ states to be a sixth level of their skill acquisition taxonomy, it is a particular iteration of performance at the two highest levels (proficiency and expertise) in which monitoring activities have stopped (D&D, 1986, 40).¹² Because monitoring has stopped and the practitioner is now thought to be completely immersed or engaged in the task and acting purely based off of fine-tuned responsiveness to their environment, Dreyfus sometimes refers to flow states as ‘mindless’ in later works.

This description of flow lines up nicely with conventional wisdom that paying attention to one’s own performance is extremely detrimental and may lead to failure or inability to perform. In other words, if the highest levels of skill are automatic, unmonitored and ‘unthinking’, then thinking about, paying attention to, or actively noticing what you are doing will lead you to be unable to perform in expert ways. It is important to note, however, that D&D suggest that no thinking at all happens during flow states, and they limit flow to those at the highest levels of performance – this sets them apart from other views on ‘flow’.

¹² D&D are criticized for their conception of skill being too passive (see Fridland, 2014) – in other words, their account is accused of suggesting that the agent is completely passive in the performance of their skill and is merely being moved by their environment – i.e. they are not acting intentionally. This might be especially evident in cases of ‘flow’ as D&D describe it. Dreyfus responds to worries of these sorts in follow-up works (Dreyfus 1997, 2002) which are covered in the next chapter, and offers an account of how this kind of action is still intentional even if passive.

For instance, another popular version of the notion of flow states is described by Mihali Csikszentmihalyi which diverges from the D&D ‘classic’ and may be less familiar to philosophers. He describes it as a state of “optimal experience”: one where a practitioner has,

“(...) a sense that one’s skills are adequate to cope with the challenges at hand, in a goal-directed, rule-bound action system that provides clear clues as to how well one is performing. Concentration is so intense that there is no attention left over to think about anything irrelevant, or to worry about problems. Self-consciousness disappears, and the sense of time becomes distorted (Csikszentmihalyi, 1990, pg 71).

For Csikszentmihalyi, what matters for being in a flow state is being appropriately occupied by the task such that it pushes the boundary of a practitioner’s current skill level but keeps out intrusive self-conscious thoughts. In order to break into this kind of concentration, practitioners often have set self-determined goals for their performance (usually, incremental performance markers) which generates sufficient intrigue and reward in a task to make it engrossing (ibid, 74). Similarly, they must cultivate an enjoyment in the rewards of the task, for example their progress and improvement of skill, etc.

For instance, a runner may set themselves the following kinds of goals; running the same course in decreasing amounts of time; running the same distance on increasingly difficult terrain, (maybe increasing inclines or running on unpaved roads, etc); running a given distance with fewer or no breaks. All of these would make their running goal-directed and would give them a marker of their performance with which to see their improvement. Trying to achieve these goals would also generate enough engagement to keep the practitioner’s attention, thereby helping them drown out other

thoughts or concerns, and make the performance rewarding, especially as they improve and gain satisfaction from this improvement.

Built into this description is the idea that, if one does become distracted from one's performance and attention is pulled elsewhere, flow will not be achievable. In some places, Csikszentmihalyi refers to this as 'psychic entropy', which he understands to be a disordering of consciousness that results from attention being pulled in various directions (ibid, 37). In other places, he brings up 'self-consciousness' as a likely indication of insufficient concentration on the task and so an inability to break into a flow. In the first sense of disordered attention, i.e. 'psychic entropy', Csikszentmihalyi seems in line with more traditional non-intellectualist thinking about flow: if one becomes distracted from the primary task, either because of fear, worry, anxiety, etc., flow will not be achievable. However, Csikszentmihalyi's understanding of self-consciousness is perhaps more nuanced than the usual suggestion that 'thinking' breaks down flow states.

During a flow state, Csikszentmihalyi would seem to suggest that practitioners are concentrating and thinking about the task at hand – they are trying to reach various self-determined goals and likely thinking about how best to achieve them. So 'thinking' alone is not the problem. 'Self-consciousness' could perhaps be better understood as focusing on oneself performing the task, rather than on the performance of the task itself. If one's thinking becomes self-referential in this way, that might be enough to introduce new pulls on attention, such as fear or anxiety, and so break a flow. In this way, there is at least a possible interpretation of Csikszentmihalyi's view that breaks away from traditional non-

intellectualist positions and perhaps foreshadows some of the intellectualist thinking to be described later in this chapter.

Furthermore, for Csikszentmihalyi, flow can happen at any level of skill from novice to expert. Again, he suggests that challenge is what causes a practitioner to concentrate and can lead to flow. As long as a practitioner is challenging their current skill level, wherever it may be, they can enter a flow state. In fact, plateauing or stagnating at a given skill level will often cause the practitioner to become bored and lose the flow state they were previously achieving. The desire to return to a flow state is what Csikszentmihalyi argues causes practitioners to continue to push themselves and improve their skill (ibid, 75). However, it seems that the degree of flow, and satisfaction in it, will increase as the complexity of performance increases. That is, even though one can experience flow at any level of skill, those with the highest levels of skill are more likely to experience the deepest and most rewarding levels of flow¹³.

¹³ Csikszentmihalyi also notes that there are three factors at play as to whether one enters a flow state (note that much of his actual writing on this front is outdated and would likely be considered by most insensitive on several fronts by today's standards – I try to present his arguments here as generously as possible). One is whether and to what extent one deems the activity meaningful. This partially explains why, to use his example, someone might find great flow in working or raising children but not in playing a game (a more traditional example of a flow-inducing activity) - said person values the former much more than the latter. This, of course, will also vary cross-culturally both in terms of what is understood to be valuable in different cultures and which cultures foster conditions that are more conducive to flow states (Csikszentmihalyi, 1990, 76-83). The second factor is the extent to which someone is able to 'restructure their consciousness' to make flow possible. People whose consciousness is too loose and unstructured (i.e. those with stimulus overinclusion or attentional disorders), as well as those with overly-structured and rigid consciousness (i.e. those who are overly self-conscious or self-centered) will have difficulty entering flow states. He also suggests that genetics and upbringing might play a role in one's ability to 'restructure' one's consciousness. For example, those who are very efficient and effective at directing their attention to key points of tasks and those who come from stable and supportive familial situations while growing up would have better psychological foundations for entering flow states (ibid, 83 – 90). However, most important, it seems, is the subject's 'non-self-conscious individualism' or the extent to which they are intrinsically motivated to do their best in all circumstances but not for the sake of advancing their own personal interests. This means that they are less likely to be deterred by external factors, are

Putting this all together, then, we can see that, even within non-intellectualist readings of flow¹⁴, there is disagreement about whether flow involves cognitive control or directives, and about whether one needs to be an expert or high-level performer in order to achieve it. This could lead us to doubt whether the criteria laid out by Christensen et al. describing the nature of automatic behavior in skillful performance are all necessary for there to be flow – indeed, if we take Csikszentmihalyi’s view seriously, we might think that they are not all markers of automaticity and flow states. Specifically, that failing to meet the standards of reduced and/or disruptive attention is not necessarily a hindrance to multi-task tolerance, reduced cognitive effort and reduced memory. This, already, might begin to suggest that there is some room for ‘thought’ and conscious processing in skillful performance, but we will see more on that as we keep going.

Regardless of the details of the conception of flow, there are some empirical studies such as (Beilock et al., 2002; Beilock & Carr, 2004; Ford et al., 2005; Gray, 2004) which seem to support the general tenets of the non-intellectualist interpretation. In each of these, expert athletes perform a skill while either directing their attention to a specific aspect of their movement or engaging in a second, unrelated task. For example, in the first condition, athletes may be asked to report, at the sound of tone, whether their bat is moving up or down and, in the second condition, asked to identify whether the sound of a tone was either high or low. All of these studies report that experts performed worse in

more able to ‘objectively’ survey their situations and observe opportunities, and so to set flow-inducing goals for themselves (ibid, 90-93).

¹⁴ Though non-intellectualists are usually the ones most concerned with ‘flow’ – its achievement and influence - Csikszentmihalyi is arguably not a non-intellectualist in his reading of flow and may be understood as a first instance of a hybrid view to aspire towards. More on this later.

the skill-focused condition than in the unrelated-task condition. Further studies (Beilock et al., 2004; Gray, 2004) also demonstrate that athletes perform worse on the skill-focused condition than in conditions with no additional task (single-task condition). All of this is taken to support the idea that attending to one's performance lessens one's ability to perform well, and so support the general non-intellectualist position.

1.3 – Intellectualist Accounts of Skill

However, the view that automaticity is central to skill has been challenged in favor of larger degrees of cognitive control or intentional action during skilled performance (see (Fridland, 2014; Montero, 2010; Papineau, 2013; Stanley & Krakauer, 2013; Sutton, 2007; Sutton et al., 2011; Toner et al., 2015, 2016). 'Cognitive control' refers to "the function of the cognitive system that allow people to regulate their behavior according to higher order goals or plans" (Verbruggen et al., 2014, p. 497) which is thought to include, "cognitive processes that are conscious and intentional in nature such as the use of rules or principles and mental representations" ((Toner et al., 2015).

Several researchers have pointed out that skilled performance often requires cognitive control for its usual performance and can also account for the accepted criteria of fluency and flexibility. For example, Fridland (2014) understands skill to be measured "by an agent's ability to respond to both expected and unpredictable environmental circumstances and to revise her strategy accordingly (...)" (2731). Performance against these measures is, according to Fridland, what defines the skillfulness of an action and requires an incredible amount of control. Fridland thinks it is control which accounts for the 'exact, nuanced' ways in which a skilled practitioner performs, modifies, revises and

guides their performance, and that “control is evident not only in the smooth, elegant execution of an uninterrupted action but in the appropriate response and recovery to variable factors as well” (2731)¹⁵. In this way, it seems Fridland understands control to be at the root of both fluency and flexibility of skill. Furthermore, she explains that she understands controlled processes to be “flexible, manipulable, subject to learning and improvement, responsive to intentional contents at the personal-level, and holistically integrated with both cognitive and motor states” (2732)¹⁶. With this understanding in mind, she argues that control should be taken as the core of skill and should therefore be at the heart of any philosophical account of skill.

In line with this line of thinking, some common experiences have been cited to support the intellectualist interpretation. Christensen et al. (2016) list them as follows:

1. skilled performance often involves *strategic focus*, in which the practitioner is attending to certain aspects of performance, especially their goals, context, methods, etc., in order to perform the task. As such, their focus is directed in such a way as to guide their actions as efficiently and effectively as possible.
2. Instances of *action slips*, in which the practitioner performs the wrong action when not paying enough attention to the larger task at hand. In this

¹⁵ Fridland goes on to clarify that she understands control to be largely *responsible* for the adjustments and guidance in skilled performance, but not *identical* to them.

¹⁶ Fridland also stipulates that, in this paper, she is not making any claims about whether control is propositional, conceptual, representational, procedural, etc. She says it may be all of none of these, but her primary aim is not to address this question. Instead, she merely wants to point out the centrality of control in skill.

way, not utilizing the appropriate cognitive control on the task at hand can lead to mistakes in performance.

3. *increased attention* in response to challenging or difficult conditions. That is, the more difficult the task, the more attention will be required in order to perform well. This might have to do with the number of variables the practitioner must keep track of, or the sheer difficulty of a single action, etc.
4. *increased cognitive effort* in response to challenging or difficult conditions. In this case, the practitioner will not only be required to pay careful attention to various aspects of performance, but they will also have to assess their situation, weigh and decide between different potential courses of action, etc. In this way, many cognitive faculties will come into play when performing well, especially at high levels.

Expanding on these ideas, several authors concretize the kinds of experiences above and aim to demonstrate some weaknesses of the non-intellectualist position. For example, (Ericsson, 2006) points out that, after a certain degree of acquisition, skills will often plateau at a given level of performance which, in order to be improved upon, will often require the practitioner to start imposing cognitive control on their behavior again. This common experience points towards a much larger degree of cognitive control in skill performance, especially within expertise, than the ‘automatic’ accounts of skill seem to allow for. For example, D&D’s view doesn’t seem to make room for plateaus in skill or for active effort to surpass them – they seem to assume that one’s skills progress in a

continuous manner and that intentional direction and rule following will gradually disappear as one reaches increasingly higher stages of expertise.

Similarly, Toner et al (2015) explain that “competitive performance regularly presents the skilled performer with situations which are challenging or relatively novel and that conscious and critical deliberation is essential if these context-specific demands are to be negotiated” (1138-9). In other words, in order to skillfully perform in competitive situations, practitioners will have to be constantly paying attention to context-specific demands and responding to them in conscious and deliberate ways. Here, too, there is push back against positions like D&D’s. The suggestion is that being appropriately responsive to an environment will not be enough to adequately perform – the more difficult the task and the higher the stakes, the more cognitive control will be required. Leaving performance at this level to an ‘unthinking’ flow leaves out many experiences of expending a great deal of cognitive effort, especially at the highest levels of performance.

1.3.1 – Thinking and ‘flow’ states

Contrary to the non-intellectualist belief in ‘flow’ and the idea that paying attention to one’s performance will be detrimental to it, Montero (2010) questions whether attention really is as harmful as is usually assumed. She begins by pointing out that many of the accounts that support the danger of attention are based on anecdotal evidence – for example, the Knoblauch case discussed above. Montero, however, addresses the same example and notes that Knoblauch himself never endorsed the

‘thinking disrupting flow’ interpretation and said that he didn’t know why he couldn’t throw anymore.

She also counters with a different anecdotal example, that of classical guitarist Tobias Schaeffer. Schaeffer claimed that even though he practiced diligently and was able to play ‘automatically’ on his own, he would often ‘black out’ and stumble during performances because he seemed to be ‘thinking about what he was doing’. However, after discussing it with his teacher, she advised him to start thinking about his movements both during practice and performance. At first, this shift in practice caused a dip in his performance. However, he was eventually able not only to return to his previous level of performance, but also improve his overall performance dramatically with no more ‘black outs’ during performance. Contrary to the non-intellectualist intuition, Schaeffer’s case seems to be one where relying on automaticity was extremely harmful to his performance and it was actually learning to pay more attention to it that turned things around.

Montero also questions the empirical data presented above that seems to support the non-intellectualist, anti-attention intuition. Montero suggests that these results (in which athletes perform worse in conditions where they must attend to a particular aspect of their skilled performance) don’t support the idea that the problem is attending to one’s performance in itself, but in attending to *particular aspects of one’s movement*, as these subjects were instructed to do. Even if it is true, Montero suggests, that paying attention to an aspect of performance like this is detrimental, it does not follow that attention to performance at large is also detrimental. And so, she suggests that the scope of these results should be appropriately reigned in.

Debunking aside, Montero suggests that the most important evidence against the non-intellectualist ‘flow’ intuition is that even experts are constantly striving to improve and that this requires active attention to performance. That is, athletes at peak performance do not simply reach the peak and cease to improve – on the contrary, even they consistently try to improve their performance either because they are trying to stay ahead of ever-evolving competition or because they simply want to keep out-doing themselves. The only way to do this, Montero suggests, is to exert active control over their training and on-line performance.

To illustrate, she considers the example of Tiger Woods who, despite being the best golfer in the world at the time, saw room for improvement in his own performance and decided to change his golf swing. While he was working on this new swing, he performed significantly worse than he had before. However, once it was mastered, he went right back up to the top of the charts. One could read Woods’ poor performance in the process of learning the new swing as an instance of ‘thought’ interfering with ‘flow’, but Montero suggests that this is better explained by the new swing still being relatively new to him rather than him ‘thinking’ about it.

1.4 – Hybrid Accounts of Skill

All in all, we can see that the non-intellectualist accounts highlight a more phenomenologically intuitive approach backed by some empirical evidence towards automatization and emphasize the responsiveness one must have towards the circumstance, one’s environment and past experience. Intellectualists, on the other hand, highlight other aspects, both in and around performance itself, to help demonstrate the

insufficiency of ‘mere’ automatization and the necessity of at least some cognitive control through both practice and performance. In this way, we can see that non-intellectualists may not be accounting for important aspects of performance highlighted by intellectualists, and that non-intellectualists may not be fully accounting for the kind of engagement and responsiveness that non-intellectualists highlight.

One solution to this debate might be to look for a hybrid view – one that maintains the intuitiveness of the non-intellectualist and intellectualist accounts respectively, but also finds a way to bridge the divide between them. One common mechanism of existing hybrid views involves delineating various levels of control of skill. For example, Christensen et al. (2019) offer the following taxonomy: The first (and highest) level of control is *strategic control*, which involves the overseeing of an extended course of action aimed at achieving a goal or completing a task. Importantly, though this level is often (but not necessarily) conscious, it does not interfere with action. Unlike cases of ‘choking’ or ‘the yips’, then, intentional guidance need not always be understood as detrimental to performance – in fact, on this account, it is necessary. Second, is *situation control*, which involves deciding what actions need to be performed in the immediate circumstance such that the goals specified in *strategic control* are achieved. Lastly, there is *implementation control*, which involves overseeing the execution of actions that have been handed down from *strategic control*.¹⁷

¹⁷ Fridland (2014) offers a similar 3-level account of control: her first level is *strategic control* which she identifies with goals, plans and strategies that are used by the agent to guide their skill across multiple instantiations. This part of skill is often, but not necessarily, conscious and are usually accessible to the agent. The second is *selective, top-down, automatic attention* in which relevant features of an environment are selected for attention based on the agent’s goals, plans and strategies. Fridland suggests that this level is automatic and usually not conscious to the agent but, importantly,

These levels of distinction open up a clear space for a series of ‘hybrid’ accounts to be put forward – in short, accounts in which skill is composed of cognitive and motor components, the former involving control while the latter involves automaticity (Beilock & Gray, 2007; Christensen et al., 2016; Papineau, 2013; Stanley & Krakauer, 2013). In varying ways, such accounts can emphasize the automaticity of skill by pointing towards ‘lower’ level of implementation control, often aligned with motor control or motor acuity. Conversely, they identify the ‘intelligent’ or intentional parts of skill with the higher-levels of control.

For example, (Papineau, 2013) argues that lower-level implementation control, which involves motor control, should be understood as automatic or reflexive for several reasons. First, because of the speed at which implementation of actions occurs, it is unlikely that cognitive control is active at this level. Here, Papineau cites studies that seem to show that cognitive processing of stimuli around oneself is happening much faster than its arrival at conscious awareness, and so deliberation or decisions about how to act would be too slow to keep up with the pace of real-world skill. Second, he points to research around the ventral and dorsal streams of the brain – the former is understood as conscious while the latter is understood to be unconscious (Goodale & Milner, 2010; A. D. Milner & Goodale, 1995, 2008). This research has shown that the unconscious dorsal

are integrated with personal level, intentional states – presumably those at the strategic level of control. The third is *motor control* in which automated motor routines that are learned through practice and training get implemented. These motor routines are automatic but the practice that leads to them involves a great deal of explicit attention and, once acquired, can be modified in various intentional ways by the practitioner. Though there is great similarity between this account and Christensen et al.’s, I personally find Christensen et al.’s formulation to be cleaner and easier to read/write, so I chose to use it.

stream is responsible for ‘online visual control of overlearned movements’, such as highly rehearsed skills, and so it is much more likely that the dorsal stream is responsible for successful execution of motor skill than the ventral stream¹⁸. Third, Papineau appeals to the manner in which phenomenon such as ‘choking’ and ‘the yips’ are caused by conscious thoughts interfering with the smooth and successful execution of skill. All of this, he suggests, stands in favor of implementation control/motor skill being automatic, unconscious and ‘unthinking’.¹⁹

For Papineau, then, automatically deployed motor routines involved in skill are separate from any intentional or cognitive control. However, intentional states, such as those in the strategic level, set parameters for motor skills/implementation control to operate within. In other words, a practitioner might set themselves the strategic intention of acting in a certain way – say, playing aggressively against a particularly pushy opponent. In addition to appropriate intention, Papineau argues that concentration, or keeping your intentions in mind, is necessary for skilled performance, especially in the maintenance of precision. For example, if, while playing soccer, a player gets distracted and starts thinking about their weekend plans, the fact that they had identified and settled

¹⁸ Similarly, building on (Butterfill & Sinigaglia, 2014), Christensen et al. (2019) consider the position that, because of the functional differences between cognitive and motor systems, the two should be understood as largely independent of each other – therefore, it is unlikely that motor systems, those active at the implementation level of control, have much influence from cognitive systems and could be understood as primarily ‘automatic’.

¹⁹ In (Papineau, 2015), Papineau puts the same thought slightly differently by first making a distinction between ‘basic actions’ (i.e. things you know how to do directly without having to decide to do anything else) and their ‘components’ (i.e. bits of behavior that make up ‘basic’ actions). He uses the example of tying one’s shoelaces as a basic skill, and of laying one lace on top of the other, making a loop in one lace, etc. as components. He goes on to argue that skilled performers need to think about *which* basic actions they will perform but that, for many of the same reasons discussed above, they cannot and do not think about *how* to perform basic actions, or in other words, their components.

on an aggressive manner of play will not be of much use since their intention is not being appropriately connected with implementation. But, regardless of what happens at the higher levels of control, implementation will happen without conscious, intentional control – for Papineau, skilled performance at this level happens in a fixed, automatic, reflex like manner.

Moreover, such accounts also highlight the degree of cognitive control that happens at the higher-levels which, they suggest, are far less likely to be automated. For example, Fridland (2014) explains that “because strategic control is concerned with the global, non-domain-specific features of a performance,” it will not become automated in the same way as lower levels (2745). Similarly, Christensen et al. (2016, 2019) focus on the degree of complexity that must be happening at higher levels of control. They point out that if skill was fully automatic, that would suggest that even the responses at the higher-levels of control are fully automated. However, automaticity takes time to develop and generally relies on repeated experience within a fixed domain of practice. Given the degree of variation that can occur in the practice of any skill, it seems extremely unlikely that the practitioner will have fully automated responses set out for the gargantuan number of possible scenarios they may encounter and/or those that they have not yet actually encountered. Because of this, a fully automated skill would likely only be able to function properly in a small number of instances. It seems unlikely, then, that automation at the higher levels of control could deliver the kind of flexibility required to perform well in anything beyond the most basic of skills.

An intuitive position based on this line of thinking could maintain some of D&D's intuitive view and suggest that cognitive control gradually shifts from lower to higher levels as the skill is learned. When a skill is first being learned, the lower implementation level is difficult and requires a good deal of attention and cognitive effort to learn to master. With time, though, this level becomes increasingly automatic and leaves room for the practitioner to focus on higher levels of situation and strategic controls. Even as far back as 1899, (Bryan & Harter) conducted a study on telegraphers that confirmed this general transition up the control-chain. They observed that novices started with learning letters, then gradually transitioned to words, then phrases and sentences, and finally to the meaning of a message. In this way, they gradually transitioned to increasingly higher degrees of automaticity at the implementation levels and could eventually come to focus their cognitive efforts on higher levels of control. Similarly, (Vallacher & Wegner, 1987) put forward an account in which the difficulty of action in initial stages of skill acquisition keeps attention at lower levels of implementation. Eventually, increased facility with the skill creates 'chunking' of actions, or larger action units, which allows for conscious focus to move higher up the control-chain of an action.

1.4.1 – 'Choking' and 'the Yips' on Hybrid views:

(Papineau, 2015) puts forward an account of cases where skilled performers fail to perform at their usual levels, which he generically calls cases of 'Not Having Your Mind Right'. He notes that 'Choking' and 'the Yips' are two subspecies of this larger

generic kind and that, though they are often run together in the literature and overlap with each other, they are nevertheless distinct phenomena.

Recall that Papineau's hybrid account mentioned above posits concentration, or keeping one's intentions in mind, as a crucial part of skilled performance. He begins with 'Choking' which he identifies as a loss of focus caused by some manner of pressure in on-line performance. For example, a performer might be expecting to do well, and even reaching the end of their performance at a very high level. But suddenly they get distracted from their actual performance and start to perform poorly.

'The Yips', on the other hand, are cases where the performer shifts their focus away from the strategic levels of performance and towards what he calls 'components', or the implementation level, of the intentionally chosen skill being performed. Papineau notes that 'the Yips' is observed in sports like golf, tennis, cricket and baseball, and only seems to "involve [those] actions initiated by the performer, rather than a response to a fast-approaching ball". In other words, when a ball is hurtling towards you, there simply isn't enough time for 'the Yips' to set in. But if you are about to throw the ball, there is "time to start worrying about the components of their basic actions and thereby undermine their performance" (305)

For Papineau, then, cases of choking might be understood as a disconnecting from higher levels of control, i.e. strategic and situational control. 'Yips', on the other hand, might be caused by inappropriate attention to lower levels of control or towards 'components' of performance. In both cases, some important break in control has happened, one 'cognitive' and one 'implementational'. But the same levels of hybrid

control could be used to understand what's happening in these cases and can pull from both intellectualist and non-intellectualist strands at the respective levels at play.

1.4.2 – Complications – Control in Implementation

However, the kind of hybrid view offered above seems to run into some issues. For example, there are those who argue that, even at the level of implementation, a great deal of cognitive control is necessary. For example, Christensen et al (2019) suggest that, for many skills, there is a great deal of complexity at the level of implementation that arises from the variability of its application across instances. In other words, each time the skill is performed will be (at least) slightly different from any other instance and will require (at least) slightly different implementation of the skill. Similarly, even a relatively simple action may require sensitivity to turbulent or changing environment, and so the practitioner's responsiveness, even at the implementation level, might require some cognitive control. Because of this, Christensen et al. argue that cognitive control must be in play to monitor and adjust implementation to ensure that the practitioner's goals are achieved.

Similarly, Fridland (2014) submits that automated processes at the level of implementation/motor control need not be understood as either fixed or 'mindless'. She cites a great deal of empirical evidence suggesting that "automatic processes are not simply processes that are executed quicker or more efficiently than those that are not automatic, but rather, that in their automation, processes undergo distinctive processing changes" (2748). In other words, automatic processes are not simply faster implementations of lists or algorithms used in non-automatic behaviors, rather a

qualitatively different computational process that continues to change even after initial automatization. As such, even ‘brute’ or low-level implementation of skills can be understood as simultaneously automated and intelligent.

Widening the scope of this conversation, Toner et al. (2015) focus not only on “on-line” performance, but on training and pre-performance rituals. Contrary to a Dreyfusian understanding of automatic or ‘habitual’ implementation, Toner et al. suggest that “winning might require not simply performing as one has in the past, but performing better than ever” (1133). They suggest that automaticity at lower levels is actually detrimental to performance as it may mean the practitioner is not ‘somaesthetically aware’ of their movement efficiency and so cannot identify disruptions to automatic movements and rectify them. Because of this they note that, even at expert levels of performance, practitioners engage in “conscious and deliberate attempts to refine and improve [their] movement proficiency” within their training regimes (1133). In other words, the most elite athletes still aim to improve their performance and so will continuously focus on altering and refining aspects of their technique in order to achieve this goal. As such, training regimes for even elite athletes may require continuous attention to the implementational level of their control so as not to plateau at a given level of performance.

With regard to pre-performance routines (i.e. instances where the practitioner has sufficient time to evaluate their current circumstances before engaging in on-line performance), the practitioner may deliberately calculate which of several options is the best in the current circumstance. They may also bring to mind similar instances of past

performance to help in practice swings, determine how much to deviate from their ‘usual’ performance, etc. These instances of deliberation may happen in the span of a few seconds and need not be long or taxing – in fact, Toner et al. suggest that, for the expert, this kind of deliberate calculation is often short and efficient.

Indeed, Toner et al. cite evidence from (Jenkins, 2007) that suggests up to 70% of the participants²⁰ used at least one cue from their pre-performance routine during their on-line performance. Indeed, there is evidence that athletes use their bodily awareness to create ‘instructional nudges’ that guide their implementation control (Gucciardi & Dimmock, 2008; Jenkins, 2007; MacPherson et al., 2008) and that the most effective of these cues are those that require the performer to be aware of the general feeling of their movement while executing a task (Mullen & Hardy, 2010)²¹.

Sutton (2007) gives a slightly more nuanced account of how this kind of verbal, instructional nudge would operate for the player, both in practice and in on-line performance. He asks the reader to recognize that “language has cognitive functions quite apart from its role in communicating or in translating our thoughts for public consumption”. He then suggests that verbalizing a ‘nudge’ could, “for example, stabiliz[e] the cognitive flow just enough to help us reorient it”. In other words, “these cognitive short-cuts can sometimes give us just that bit of extra flexible independence to

²⁰ Who were a group of 113 European tour gold players.

²¹ Sutton et al. (2011), whom we’ll see more of later, describe this kind of process (of using a verbal cue from practice to guidance on-line performance) in much more embodied terminology. They describe it as one in which the performer, “is using these verbal components of multi-modal embodied routines to distribute intelligence, coordinating or often re-setting and rechunking patterns of movement or affect or mood, as one among many forms of scaffolding that support the embodied rebuilding of action sequences from the inside” (93).

adjust the way we respond”. So, rather than understanding these instructional nudges as completely top-down instructions, we should instead understand these cues as “material symbols with temporary but crucial causal roles as ‘a new fulcrum for the control of action’” such that “a complex bodily pattern or set of possible movements can be compressed into and partly cued by a phrase or memory or ingrained image, bringing the payer back to, rather than away from, the well-learned habits” (774).

All of this seems to suggest that practitioners often consciously adapt their on-line performance at the implementation level using cues and routines built up in training and pre-performance and so exert some ‘cognitive control’ over this level, too.

1.4.3 – Complications - Situational Control:

A further complication for many hybrid views is that some research suggests that even largely automated processes in situational control should be understood as intelligent (Fridland, 2014; Pylyshyn, 2003; Wu, 2014). Citing evidence from Pylyshyn (2003) and Wu (2014), Fridland (2014) explains how situational control involves a great deal of selective attention, which is sometimes thought to be one of the defining features of expertise. That is, attention is trained over time and with intentional practice such that the expert quickly and automatically notices many salient features of their environment and will pay attention to specific parts of a scene. However, this attention is also importantly in tune with the practitioner’s strategic goals – what the practitioner pays attention to lines up neatly with what they’re trying to do in that moment. As such, this level of control is both automatic and mindful. So again, the tidy hybrid position whereby

the domain of control is in the highest levels and automaticity in the lowest level may turn out to be too simplistic.

1.4.4 – Exemplary Hybrid View – Sutton et al.’s AIR

How, then, are we to settle this debate? As the dissertation continues, we will return to issues of automaticity and control as they relate to memory of skill and look back at this debate in a different context. As such, I will not yet offer my own view or understanding of the topic. However, in order to move on, it will be helpful to have a concrete example to hold on to and give us a sense of what I think the right sort of relationship is. With this in mind, I present here (Sutton et al., 2011)’s Applying Intelligence to the Reflexes (AIR) model²².

Sutton et al.’s view suggests that, rather than making skill fit to the form of the dichotomous debate (intelligent or ‘mindless’), what is key for skill is the manner in which attention is utilized in performance. Skilled performance will sometimes require attention to one’s body unfolding in action, to memory and past experience, to semantic thought, to context around oneself, etc. – attention must be able to flow to all of these things freely and incorporate them into performance in real time. They come to this conclusion by noticing several things. First, they note that “genuine expertise often requires the rapid switching of modes and styles *within* the performance context” which means that action “must thus be open, under certain circumstances, to the influence of explicit knowledge, specific memories, or particular decisions” (original emphasis, 93).

²² The AIR account is also developed in other papers (Geeves et al. 2014) and (Geeves et al., 2008) . Sutton, himself, has described similar/compatible positions in other works such as (Sutton, 2007).

In other words, skilled performance will require the practitioner to pay attention to a great number of things and be able to switch between them – this includes their personal experiences that have led to their skillful abilities but also to various context dependent factors of a given instance of performance.

In order to be able to move their attention around in this manner, Sutton et al. argue that skilled performers have “opened their ‘reflexes’ up into acquired adoptive patterns, and constructed (...) dynamic repertoires of potential action sequences which can be accessed, redeployed, and transformed appropriately” (ibid, pg 96). In other words, a practitioner’s ‘reflexes’ or responsiveness to the situation must be trained such that they can flexibly move between learned ‘patterns’ or ‘repertoires’ of behavior and the ability to adjust or deploy those patterns as needed in response to real-time circumstances. As such, Sutton et al. understand expert skill to be “constituted by a dynamic responding that involves retrieval of the most effective (combination of) learned material given the unpredictable, contingent contextual demands with which a [skilled practitioner] is faced during performance, and the integration with and expansion of this material in line with these demands” (Geeves et al. 2014, 9).

All of this means that, for Sutton et al., skill “relies on a mindedness that facilitates the dynamic flexibility of attention, allowing it to be allocated freely and in a way that best meets contingent contextual demands” (Geeves et al., 2014, 3). With this kind of flexible attention, experts are able to monitor their performance and make on-the-fly decisions throughout. Because these decisions may involve drawing on past memories to plan and strategize as well as somaesthetic evaluations about how to modify motor

performance, this means that, even in the most deeply ingrained skills, we “retain significant levels of care, attention, and kinetic awareness” (Sutton et al, 2011, p88). In this way, skillful performance is (perhaps ambiguously) infused with control and automaticity throughout. In fact, Sutton et al. argue that it is precisely because experts are so used to engaging in this ‘minded’ manner that they are able to so flexibly redirect their attention. This is also what allows them to reconstruct ‘weaker’ versions of behavior during practice and transform behavior during on-line performance.

Sutton et al. see themselves as offering a view between intellectualist and non-intellectualist positions that maintains their respective intuitions and strengths. For example, they see their view as being in-line with intellectualist understandings in which cognitive control, strategy and the like are developed and carefully implemented throughout performance – indeed, at every level of the three-tiered control hierarchy described previously, including implementation. Note, for instance, that they maintain the significance of planning and strategy in response to the practitioner’s goals and other context-dependent circumstances – in other words, they highlight control at the strategic and situational levels of control. They also highlight the role of constant monitoring of the implementation level and the mechanisms through which even this ‘low’ level can be impacted by cognitive control

Similarly, Sutton et al. understand their view to be maintaining some of the Dreyfusian non-intellectualist intuitions. They maintain the role of a flexible responsiveness to an environment or circumstance in a way that relies heavily on past experience and ‘automated’ pattern recognition, etc. Much of what Dreyfus would say

about intuitive highlighting of salient features of a scene, being engaged in holistic patterns of behavior, etc. can all be maintained on the AIR model. Sutton et al. understand Dreyfus to be primarily underscoring the depth of involvement in a task such that a practitioner is ‘no more aware of it than [they are] of [their] own body’. Their view, they suggest, is not in conflict with this since “knowledge processes underlying expertise may well be unconscious, but the AIR approach believes that they are accessible and articulable in the right conditions” (Geeves et al. 2014, 8). Indeed, their view seems to allow a reading where the mindedness that allows for flexible attention is itself an achievement of a high level of engagement in a situation and perhaps of a certain kind of flow. Sutton et al. merely suggest that, when the practitioner is within the scope of these conditions, they can assess performance-related information relative to their situation and follow this up with appropriate, on-the-fly modifications. This means that the AIR model allows for automaticity in all three levels of the control hierarchy, as well. Of course, implementation and situational control may be relying on learned patterns and ‘automatically’ deployed repertoires, etc. However, strategic control at the highest levels may also ultimately result from the right kind of familiarity and experience with said performance and so involves a kind of flow or automaticity, as well.

Learning from this kind of approach, I suggest a revision of our second criterion:

Criterion 2*: skill is typically manifested in fluent and flexible action. The mechanisms underpinning fluency and flexibility will utilize both automaticity and cognitive control throughout performance in ways that are perhaps ambiguous and difficult to separate.

As such, rather than focusing on when automaticity or cognitive control is at play, focusing on a third feature, perhaps how attention is utilized in performance, may be a better way to account for fluency and flexibility in skill.

1.5 – Habit vs Skill:

Recall that procedural memory, the ultimate focus of this dissertation, is often understood as the domain of skill and habit. Having gone through all of the above on skill, it may now be helpful to more carefully examine its relationship with its close relative, habit. The classic treatment of the difference between skill and habit comes from Gilbert Ryle (Ryle, 1949). Ryle understands habit as being a continuously repeated behavior that does not change between practices: it is acquired through ‘drill or conditioning’, it is rote, and a ‘mere replica’ of previous instances of the behavior. That is, habit is the kind of behavior one repeats automatically, without much thought or deliberation about doing so. Furthermore, because habitual behaviors are continuously reproduced and performed in the same ways and are not performed with any or much self-monitoring, there is no drive to improve the performance of the behavior. Once a habit is set, it’s performance will likely stay that way indefinitely. This leads Ryle to say that habits are single-track dispositions: dispositions to act in one particular way under the right circumstances.

Skill, on the other hand, falls under the category of ‘Knowing-How’, which Ryle understands to be an intelligent way of acting (which need not be guided by any conscious thought or deliberation). Skill is gained through training, which Ryle distinguishes from drill. Training involves self-monitoring and critique – while training,

one is trying to improve one's performance and so must be applying awareness and judgement throughout. This is what makes skill 'intelligent' - it involves a sensitivity to, and awareness of, one's circumstances driven by the aim of improvement. Because of this, Ryle understands skill to be dispositions that are 'indefinitely heterogenous' – when one has a skill, one has a disposition to behave in any number of ways in accordance with one's context. Along the same lines, skill is practiced for improvement – one must consistently monitor one's progress in order to identify mistakes and find ways to avoid them or improve one's performance. Skill, then, is also 'heterogenous' in that it modifies previous performances (Ryle, 1949, 42-45).

Even so, Ryle highlights something that has, so far, been implicit in our conversation; namely, that skill is aimed at 'getting it right' and that the intentionality of the practitioner is key. Indeed, for him, this is the key distinction between skill and habit. Our conversation up till now has relied on similar notions, especially within the intellectualist camp, to push for an understanding of cognitive control as crucial to skill. In this way, Ryle's view on this distinction leans him closer to the intellectualist camp²³ and is something worth highlighting in our discussion. As such, I suggest the following criteria be added to our list:

Criterion 3: skill is aimed at 'getting it right' and is underscored by the practitioner's intentionality

²³ However, as we proceed to Chp. 2, we will see that his understanding of skill, especially as it relates to knowing-how, is more nuanced than it seems here, and highlights some aspects of what might be understood as non-intlectualist understandings. Indeed, Ryle is usually known as a non-intellectualist, in general.

Some, however, have pushed back on Ryle's strong delineation between habit and skill, the former 'dumb' and the latter 'intelligent'. (Pollard, 2006), for example, suggests a slightly more sophisticated version of habit that might line it up more closely with skill. Pollard suggests that habits are central to agency. So much so that, in some circumstances, identifying an action as a habit is an effective explanation of that action. Indeed, unlike Ryle, Pollard understands habitual actions as being more open to influence than, say, addictions or compulsions. Though habits do not require "any preceding deliberation" they are still unlike mere reflexes, or mere repetition, since the agent retains some direct power to intervene and control them. In this sense, then, habits have more in common with skills on Pollard's picture than they do on Ryle's.

Similarly, Sutton et al. (2011), for example, suggest that "habits can be flexible and adaptive as well as idiosyncratic", and so should be seen, at least in certain contexts, "as more like immersed embodied skills" (Sutton et al. 2011, 81). They suggest that our habits, such as driving in particular ways or cooking in our 'normal' ways, can be genuinely context-sensitive when, for example, we're driving under changing conditions or cooking for a special occasion. Similarly, they suggest that even more mundane habits, such as brushing one's teeth or collecting one's keys and belongings before leaving home, can "remain more or less open and responsive to any peculiarities of today's unique constellation of moods and events" (ibid, 80)²⁴. Even if the initiation of these

²⁴ In this piece, Sutton also argues for an interpretation of Descartes that paints a similar picture – in other words, that habit, in Descartes mind, "are grounded in dispositions, which in turn are grounding in the complex dynamical arrangements of physical parts". As such, our past experience is 'sedimented in associations' which explains both our 'habits' and our 'skills' – both are part of our dispositions to act in accordance with internalized past experience.

behaviors no longer requires attention, Sutton et al. argue that they still “intrinsically involve certain kinds of monitoring” and “often retain significant levels of care, attention, and kinetic awareness” (ibid, 88).

We can see that, though habits have been historically understood as merely repetitive and rote behavior, they can be understood to utilize a similar kind of control to skills. Perhaps more intuitively, habits have been understood to have a kind of automaticity and ‘unthinkingness’. Insofar, then, that habits share with skills a manner of control, we may also be inclined to transpose some of the thinking around responsiveness and automaticity of skills back into habits. The kind of automaticity and responsiveness found in habits may not be as ‘dumb’ as Ryle suggested and may look closer to the kind of ‘flow’ described by non-intellectualists.

Consider our established criteria for skill so far: namely that skill is learned, that it exhibits a complex relationship between automaticity and control in the display of fluency and flexibility, and that it is aimed at ‘getting it right’ and so underscored by intentionality. Habits can be understood to exhibit these characteristics as well. To begin with, they are certainly learned, not innate. Similarly, insofar as habits rely on past experience for their formation, are responsive to circumstance and can be monitored and regulated accordingly, they may demonstrate a fluency and flexibility in ways akin to skill. Furthermore, even a mundane habit, like brushing one’s teeth, could be understood as being formed in such a way as to ‘get it right’ and so underscored by intentionality. For our purposes, then, we can understand the relationship between skill and habit as a

very close one – one where their characteristics overlap and, in some cases, may be difficult to distinguish.

Even so, there seems to be good reason to maintain a conceptual distinction between them, despite this (potentially very significant overlap). For instance, suppose I have a habit of playing all the New York Times Games in the app every morning as I drink my coffee or tea: Wordle, the mini crossword, Spelling Bee, and the like. It seems plausible to think that all of these would utilize some skills, but the fact that I do so every morning does not seem like a skill in itself. Similarly, suppose I have a habit of playing a particular chord on a guitar with finger A rather than fingering B – I can, of course, play both fingerings and the tendency is not a question of picking the easier one or what have you – it's just a tendency. Playing the chord in fingering A might require a level of skill, but having the habit of doing so does not clearly seem to be a skill in itself, either.

So, what's the difference here? What we've covered above has pushed us to consider that the acquisition and oversight of a habit need not be understood as rote or 'dumb' as initially thought but, in both the cases we considered, there is still some level of repetitiveness and predictability that does not seem to rely on being skilled in any way. One way of understanding the difference is that habit acts as a sort of baseline from which skill can operate and elevate behavior. IN a sense, we could think of habit as an initial guide for behavior, but one that requires oversight to carry out. Because the initial guide requires oversight, there is room for more skilled behavior to come into play. Though both are based on acquisition over past experience and a sort of automatization, and both require a degree of oversight to carry out, habit is still barer than skill – it is not

as rich, refined and contextually dependent. So, rather than seeing habit and skill as completely opposed, we seem well placed to see them as continuous but distinct.

1.6 - Conclusion

By examining a debate between intellectualists and non-intellectualists in philosophy of mind, cognitive science, and psychology, we have identified two criteria of skill: namely that skills are learned and that they exhibit fluency and flexibility through a complex relationship between automaticity and control. We have also highlighted an exemplary hybrid view that provides a framework with which to proceed into later parts of the dissertation. In the next chapter, we will examine the epistemological literature around skill and, from it, extract a few more criteria.

Chapter 2

Epistemological accounts of skill:

2.0: Introduction:

Memory is usually thought to be the subject's retention of some kind of knowledge or information for future use. In so far as we are interested in an account of procedural memory, or memory of skill, an understanding of what kind of knowledge or information is at play in this memory seems critical. Indeed, the competing views on the nature of memory, in general, have often incorporated epistemological literature to bolster or supplement the author's favored conceptual framework about memory. The case of procedural memory is no exception.

As such, this chapter will be aimed at exploring the epistemological literature on skill and extracting a few more criteria for our list about the relationship between skill and knowledge. According to (Pavese, 2016), epistemology has only recently taken an interest in skill due to the rise of two prominent debates – the first being the debate of the nature of know-how and the second the debate on virtue epistemology. In both debates lies a similar dichotomy – one of trying to understand whether knowledge should be understood as a type of/exercise of skill or whether skill should be understood as a type of/exercise of knowledge. Once we consider these debates, we will be in a position to examine more general positions on the relation between skill and knowledge and utilize these understandings in later discussion of memory.

2.1 – Knowing-How vs. Knowing-That:

2.1.1 – Know-How and Skill:

To begin with, let us examine the question of what kind of knowledge is thought to be involved in skill. Traditionally, skill has been understood to involve ‘knowledge-how’, or knowing how to do something, as opposed to ‘knowledge-that’, or knowing that something is the case, knowing some facts, etc. Though there are likely many reasons for this, (Pavese, 2016) identifies three as central. First, that it is hard not to talk about know-how when talking about skill – in fact, many languages, including French and Italian, don’t have a distinct word for ‘skill’ that is separable from the words for ‘ability’ or ‘talent’. Secondly, she points out that “it is intuitive that one cannot be skilled at [something] without knowing how to perform it”, or, in other words, that being skilled at ϕ -ing seems to require knowing how to ϕ (646). For example, it doesn’t seem that one can say they are a skilled cook unless they also know how to cook. Lastly, it also seems that know-how entails skill – it’s hard to say that someone knows how to cook if they do not also actually have the skill of cooking.

However, Pavese considers a worry to this neat and tidy picture. She points out that someone could know-how to do something, for example, baking a cake, but not be particularly skilled at it. That is, they might genuinely utilize some know-how to bake a cake – they know how to make the dough, the icing, how to assemble it, etc., but the resulting cake is not particularly good. This would suggest that knowledge-how and skill can come apart. Even so, Pavese points out that the notion of ‘skill’ is a gradable one, one that admits of degrees *above a certain threshold*. So, when we say that someone is skilled

at something, what we mean is that someone's skill surpasses some contextually fixed threshold. As such, someone could be understood as knowing-how to bake a cake but not be understood as skilled at it, since their skill could be understood to fall below the given threshold. But, if we say that someone knows-how to bake a cake *sufficiently well*, that does entail being skilled at it. So, know-how *above a certain threshold* does entail skill.

Pausing for a brief interlude, let us consider this general notion. Though the above claims are specifically aimed at the relation between KH and skill, others have made similar claims. For example, Stanley and Krakauer (2013) make a similar point, i.e. that having a skill requires "being trained past a baseline" and use this as one of their two main claims about the nature of skill (we will come back to this piece later on in this chapter). This is a new idea to our existing discussion since, in the previous chapter, skill was either implicitly understood to include what might be understood as states 'below the threshold' (i.e. Dreyfus), or the only skills discussed were those of experts. As such, I'd suggest the following criterion be added to our list, namely:

Criterion 4: to be understood as skilled entails performance above a context specific threshold.

Coming back to the discussion of KH (specifically, that KH above given threshold entails skill), we can transfer whatever is said about know-how to our understanding of skill itself. The notion that skill and know-how are closely related is generally well established. However, this raises the following questions: what is the nature of KH?; and what is the relationship between 'knowing-how' and 'knowing-that', (and so, by extension, between skill and 'knowing that'). To address these questions, we

will start by examining Ryle's classic account of the distinction between knowing-how and knowing-that followed by a well-known opposing view found in Stanley and Williamson's work.

2.1.2 –Non-Intellectualist Positions:

Gilbert Ryle (Ryle, 1949) seems to have originated the distinction between knowledge-how and knowledge-that and understands them to be different modes in which intelligent behavior can be exhibited. For Ryle, 'intelligence' has little to do with what knowledge or truths one has amassed (i.e. KT) but with "the ability, or inability, to do certain sorts of things" (i.e. KH) (ibid, 27). What matters for intelligence is someone's 'competence' in using various kinds of knowledge; their ability to organize it and exploit it. In this way, intelligence seems to be equated with knowledge-how, for Ryle: one must know-how to utilize various bit of information, truths, facts, circumstances, etc. in order to be understood as intelligent, and doing this depends on having the rights sorts of skills in place.

In elaborating on the nature of knowledge-how, Ryle begins by saying that, when one says that someone knows how to do something, "part of what is meant is that, when they perform these operations, they tend to perform them well, i.e., correctly or efficiently or skillfully" (ibid, 28). However, simply meeting certain standards of performance isn't enough for knowledge-how. Ryle goes on to say that "[a] person's performance is described as careful or skillful, if in his operations he is ready to detect and correct lapses, to repeat and improve upon success, to profit from the examples of others and so forth" (ibid, 28-29). Knowledge-how, then, is defined in terms of skillful

performance that aims to ‘get things right’, and this Ryle understands to be ‘intelligent’ behavior. This means that know-how or skill is intelligent in its own right and is distinct from knowledge-that.

Ryle then considers an understanding of the relationship between KH and KT such that KT underpins KH. This view expresses the idea that “an action exhibits intelligence, if and only if, the agent is thinking what he is doing while he is doing it, and thinking what he is doing in such a manner that he would not do the action so well if he were not thinking what he is doing” (ibid, 29). He suggests that this common view is thought to support an intellectualist perspective in which KH reduces to KT and suggests that “intelligent performance involves the observance of rules, or the application of criteria” (ibid, 29). On this intellectualist reading, as Ryle sees it, intelligent performance “is to do a bit of theory and then to do a bit of practice” (ibid, 29). Ryle concedes that one certainly does usually reflect before action so that they might act ‘properly’, and that this might happen extremely quickly and unconsciously on the part of the agent. But ultimately, Ryle aims to reject this intellectualist line of thinking and demonstrate its insufficiency.

To begin with, Ryle says that “there are many classes of performances in which intelligence is displayed, but the rules or criteria of which are unformulated” (ibid, 30). In other words, many people act ‘intelligently’ or skillfully but do not have any rules or formulas available that they could share or say they follow. This would seem to suggest that KT is not directly in play during skillful performance as the intellectualist suggests. The example Ryle uses is that of a wit – someone might be very adept at making and

appreciating jokes, and might even be good at identifying bad jokes, but cannot tell anyone else what ‘recipes’ they use to do this. In fact, Ryle suggests that “efficient practice precedes the theory of it”, or that one must already have a skill in order to be able to identify and pick out points for theorizing about it (ibid, 30)²⁵. In this way, knowledge-that isn’t necessary for knowledge-how and is even secondary to knowledge-how.

However, Ryle’s ‘crucial’ objection to the intellectualist position is this: “the ability to apply rules is the product of practice” (ibid, 42). In other words, considering propositions is, itself, an act that can be done more or less intelligently, more or less skillfully. To properly use any maxims or rules that one is considering, one must use ‘good sense’ which “cannot itself be a product of the intellectual acknowledgement of any general principle” (ibid, 31). Instead, this good sense must come from knowledge-how: from skill gained from experience²⁶. Furthermore, if it was true, as the intellectualist suggests, that in order to act skillfully one needs to ‘do a bit of theory’ first, then “it would be a logical impossibility for anyone ever to break into the circle” (ibid, 30). In

²⁵ (Dickie, 2012) makes a similar point – she argues that KT requires a conceptual representation of the content at hand while KH does not. She argues that “The agent of a skilled ϕ -ing-in- w stands in a non-conceptual representing relation to w : s/he is sensitive to the difference between w and w^* , but need have no grasp of w ’s criterion of identity” (745).

²⁶ Arguments making similar points are not uncommon. For instance, (Fridland, 2013, 2014) makes a similar point – she argues that mere possession of information/knowledge (i.e. knowledge-that) isn’t enough for know-how. There must also be intelligent selection of means to ends for know-how. In even broader terms, John Searle’s notion of a ‘background’ against which one acts seems to suggest something similar – that is, that rules and propositions (or KT) don’t come self-interpreted. In other words, that rules get interpreted in or against the subject’s ‘background’ understanding of their world which seems to rely on a non-propositional sense of competences. Whatever rules or propositions one encounters will never be never offer exhaustive instruction and will always require some competence to interpret (see (Rust, 2009) for more).

other words, the intellectualist view creates a circularity: it puts ‘theory’ first in order to understand skill, but ‘theory’ itself requires skill. For example, the subject must be able to identify which bits of ‘theory’ are most relevant to their situation; they must “conside[r] what is pertinent and disregar[d] what is inappropriate” (ibid, 31). In light of this, Ryle argues that trying to maintain the intellectualist position would entail a regress: the subject must then reflect on how best to reflect, or already have knowledge-how in order to know-how to do something, and so there could never really be any manifestation of knowing-how. This regress leads Ryle to think that an account of intelligence cannot be one that requires reflection on ‘theory’, or knowledge-that.

With this in hand, Ryle offers his positive account of knowing-how. He directs the reader to two case studies in which a young player learns chess: in the first, they are taught various rules of play which they rely on heavily during earlier phases of their training, but gradually, as their expertise increases, they stop using them and may even forget them. The second case is also of a young chess player who is never taught any rules but gradually learns to play the game simply through watching others. In both cases, the young chess player has either forgotten or was never taught any explicit rules of play, decision making, etc., cannot articulate them, and yet learns to play the game at a high level and can be shown to have some understanding of these rules or maxims regardless. Ryle uses these cases to demonstrate that “we learn *how* by practice, schooled indeed by criticism and example, but often quite unaided by any lessons in the theory” (ibid, 41, original emphasis). What demonstrates someone’s knowledge-how, then, is not

propositions about the activity or rules or maxims they know about it, but their actual performance which confirms an understanding of said rules or maxims.

It should be noted that, though he is usually taken to represent the non-intellectualist view on KH, Ryle's view actually straddles some tendencies between both intellectualist and non-intellectualist views. Again, he is usually understood to be a non-intellectualist in so far as he rejects the notion that propositional or declarative knowledge, or what he calls 'theory', is necessary for skilled behavior. However, his view also incorporates some intellectualist intuitions by suggesting that skill requires monitoring and adjustment of performance guided by goals and situational feedback; in other words, as he says, skill must be 'intelligent'. As such, he follows some of the tendencies of the AIR model (highlighted in the previous chapter as an exemplary hybrid view), in which there is an important role for cognitive control in skill (especially in its use of theory and critical on-line thinking), but also for 'automated' training and experience to stand on its own and, in fact, underlie and guide whatever 'theory' is used. In this way, Ryle could be understood as another key example or model for of the kind of hybrid view we have been pushing for - a kind of blending of views that shows how both KH/KT are relevant to skillful performance and takes KH and its guiding role seriously. However, in the end Ryle settles on KH, or skill based in experience, as the more fundamental of the two, which might separate him from something like the AIR model and help entrench a distinction between KH and KT that may prove troublesome.

2.1.3 – Intellectualist Positions:

Despite Ryle's argument that KH cannot be reduced to KT, others have argued for precisely this view. Perhaps the paradigmatic statement of this position comes from (Stanley & Williamson, 2001) in which they argue that KH should be understood as propositional knowledge (in other words, as KT). To do this, Stanley and Williamson (henceforth S&W) begin by considering Ryle's own account and his denial that KH could be understood as (a form of) KT. They say that it is "quite clear that Ryle took his central arguments against the thesis that knowledge-how is a species of knowledge-that to hinge on an accusation of vicious regress" (ibid, 413) and provide the following breakdown of Ryle's position.

"Ryle's argument has two premises:

- (1) If one *F*s, one employs knowledge how to *F*.
 - (2) If one employs knowledge that *p*, one contemplates the proposition that *p*."
- (ibid, 413)

As S&W understand him, Ryle's move is to suggest that, "if premise (1) and (2) are true, then if knowledge-how is a species of knowledge-that, doing anything would require contemplating an infinite number of propositions of ever-increasing complexity" (ibid, 414). However, S&W suggest that Ryle's argument cannot get off the ground because it is unsound – that is, "there is no uniform reading of the two premises in Ryle's argument on which both are true" and so the internal structure of Ryle's argument is flawed (ibid, 416).

To establish this claim, they begin by examining premise (1) and noting that it can only apply to intentional actions. For example, when one is digesting, it is clearly false that one is employing knowledge of how to digest in doing so, because this seems to be a

case in which one is not intentionally acting and so knowledge is not directly relevant. With this clarification, S&W move on to considering premise (2). They note that Ryle assumes that the use of knowledge-that requires ‘distinct acts of contemplating propositions’, which also need not be the case – one could use their knowledge-that without formulating or contemplating any propositions²⁷. This being the case, S&W suggest that premise (2) is patently false. Though they argue that this alone suffices to undermine Ryle’s view, S&W consider an attempt to rescue premise (2) by suggesting that there could be a sense of ‘contemplating’ which doesn’t require intentionality or maybe only refers to some deflationary sense of action. If this route is taken, however, S&W point out that this would clearly put premises (1) and (2) in tension with each other and would mean that Ryle’s argument isn’t sound. In short, S&W argue that either premise (2) is false, and Ryle’s argument fails, or premises (1) and (2) are incompatible with each other and Ryle’s argument fails again.

From here, S&W move onto their own positive account of KH which is based largely in linguistic arguments about syntactic and semantic theory, though I will focus primarily on the latter. After reviewing a fair bit of literature on semantic theory, S&W put forward their own positive account of the semantics at play in ascriptions of

²⁷ Here it is worth noting that, though this seems to be the very claim they are trying to establish, they simply assert it with virtually no argument. However, in more recent work, (Stanley & Williamson, 2017) point out that the knowledge involved in skill (which includes knowledge-how) is not ‘discursive’. In other words, “It cannot be explained in non-indexical terms”. Even so, they say that such knowledge “has propositional content, nevertheless. Not all propositional knowledge is discursive knowledge” (ibid, 715). They do this to try to allow themselves to maintain the kind of un-reflective automaticity often associated with the fluency and flexibility of skill. Even so, it’s still not clear that this proves the point that propositions can be entertained without formulation or contemplation. For a more in depth version of this objection, see (Noe, 2005)

knowledge-how. S&W suggest that knowing-how statements are statements with embedded questions. For example, saying “Hannah knows how to ride a bike” embeds the question “what is the way to ride a bike?” and implies that Hannah knows the answer to this²⁸. As such, there are certain semantic implications for the sentence “Hannah knows how to ride a bicycle”. For example, it is true “if and only if, for some contextually relevant way w which is a way for Hannah to ride a bicycle, Hannah knows that w is a way for her to ride a bicycle. Thus, to say that someone knows how to F is always to ascribe to them knowledge-that” (ibid, 426). In other words, Hannah knows how to ride a bike if, given her circumstances, there actually is some way for her to do so and she can identify it as such. To do this, she must stand in a knowledge-that relation to a Russellian proposition²⁹ (in this case, w is a way to ride a bike), which amounts to knowledge-that she must work through the ordered sequence of properties and objects contained in the proposition. In short, Hannah’s KH to ride a bike consists in her working through a series of propositions about whether and how she can do so, i.e. KT. It is worth noting that S&W do not say more about what this ‘working through’ would amount to.

²⁸ In more recent work, S&W have expanded on this idea to include other types of knowledge along with KH in the embedded questions. In (Stanley, 2011a; Stanley & Williamson, 2017), they elaborate on the relation between statements of knowing-how and statements of what they call ‘knowing-wh’ (which includes knowing-where, -when, -why, -whether, etc.). In many cases, attributions of knowing-how amount to attributions of knowing-where, -when, -whether, etc. For example, knowing how to play basketball may involve knowing-when to make a drive for the basket, or knowing how to ride a bike may involve knowing whether to use the brakes as one approaches a sharp turn. As such, S&W suggest that whatever is said of knowing-how attributions should be said of ‘knowing-wh’ attributions as well, since they are fundamentally complimentary. In fact, they consider knowing-how statements to be held within the broader category of ‘knowing-wh’ statements and say that in all of them “one is in such a state in virtue of *knowing the answer to a question*” (S&W 2017, 715).

²⁹ A proposition in which properties and objects are ordered sequences

Crucially, once they incorporate theories of propositional attitudes, S&W amend their view to include modes of presentation. They suggest that there must be such a thing as a ‘practical mode of presentation’ and that this mode would be the one used in ascriptions of knowledge-how. In short, they suggest this mode helps demonstrate the difference between, for example, being told how to do something and doing it oneself. If one is simply told that “*w* is a way to ride a bike”, there is a sense in which one then knows how to ride a bike without having any *practical* knowledge of how to do so. Representing the same proposition *practically* is what is meant to make this difference here.

So, we end up with S&W’s final statement of knowing-how: the statement,

“Hannah knows how to ride a bicycle” is true “relative to a context *c* if and only if there is some contextually relevant way *w* such that Hannah stands in the knowledge-that relation to the Russellian proposition that *w* is a way for Hannah to ride a bicycle, and Hannah entertains this proposition under a practical mode of presentation” (430).

2.1.4 – Takeaway from KH-KT Debate

S&W’s view is a standard intellectualist view that understands KH to be a sub-species of KT and suggests that instances of KH amount to working through a series of propositions about how to get a task done. Crucially, however, in order to make their case they seem to make at least one vital concession to non-intellectualist views; specifically, the necessity of propositions being understood through a ‘practical mode of presentation’. S&W don’t present this as a concession and offer virtually nothing in its defense. All they say is that it is intricately tied to a complex of dispositions. S&W rely, instead, on the fact that this mode of presentation is obscure in the same way as others are and suggest that

this doesn't provide any reason to think that they're not propositional. Even in later works, Stanley maintains this general position (Stanley, 2011b). S&W have been criticized for the mysterious nature of this 'practical mode of presentation', especially since it is so central to their view (see for example, Noe (2005)). As such, the inclusion of this 'practical mode of presentation' seems to open a back door for the kind of argument Ryle originally made: one where some practical knowledge or ability (KH) is necessary in order to utilize one's KT. If this is the case, S&W's view may not prove the point they are aiming at and might instead fit in rather neatly into our larger picture.

All in all, views like S&W can be understood as highlighting the kinds of KT that would be necessary for skill, while also seemingly inadvertently relying on notions that shore up the non-intellectualist position. However, even non-intellectualist views like Ryle's make room for a great deal of 'intelligence' and active consideration of KT throughout performance. Whether KH or KT is more fundamental to skill is perhaps not the most relevant question for our purposes. Instead, as we saw in the previous chapter with automaticity and cognitive control, it may be that we should be looking for accounts that utilize both types of knowledge throughout skillful action and demonstrate that neither is more fundamental than the other. As such, we can tentatively set out our fifth criterion for skill:

Criterion 5: Skillful performance relies on an interdependence between KH and KT. More specifically, KH utilizes elements of control and KT in performance, while KT relies on at least some KH to be able to use, interpret or execute it.

2.2: Skill and Knowledge:

The debate between KH and KT is only one part of the larger debate about the relation between knowledge and skill. The wider terrain involves the delineation between procedural and declarative knowledge (which, as we will see, relates directly to the discussion of different types of memory along the same lines). Declarative knowledge has been traditionally thought to include KT, and/or semantic, factual, or propositional knowledge, and is thought to be knowledge that can be recollected at will. Procedural knowledge, on the other hand, is thought to be knowledge gained through experience and expressed through changes in behavior. It is either thought to include KH or is simply directly equated with it.

One of the primary findings that helped substantiate this distinction came from the case of HM. HM was a patient with severe epilepsy who underwent a bilateral temporal lobectomy to help treat it. Afterwards, he was found to have acute amnesia such that he would forget events quickly after they happened. However, HM was shown in experiments to have been able to improve his ability at a motor task³⁰ over the course of several days. Of course, given his amnesia, HM did not remember having encountered the task before and never felt any familiarity with it. Nevertheless, his performance improved (B. Milner, 1962). This result has been traditionally interpreted as a demonstration of the separation between declarative memory/knowledge and procedural

³⁰ specifically, tracing the outline of a star with a pencil through a mirror with vision of his own arm obscured

memory/knowledge since HM retained his procedural knowledge but not his declarative knowledge.

The lines drawn between intellectualist and non-intellectualist positions in this debate will likely be familiar at this point. Following the results of HM's case, non-intellectualists have traditionally interpreted procedural knowledge as the relevant sort of knowledge for skill. As such, non-intellectualists have sought to offer models or theories of skill that do not rely on declarative knowledge and can flesh out an understanding of what procedural knowledge might amount to. However, intellectualists have made arguments highlighting the ways in which declarative and/or propositional knowledge are necessary for skill and ways in which procedural knowledge can be understood in the same terms as declarative knowledge. Intellectualists have often done this by arguing for the necessity of representations in cognition, that procedural knowledge can be understood propositionally, etc.

2.2.1 – A non-intellectualist account of skill and knowledge

For one of the most well-known accounts of a non-intellectualist relation between skill and knowledge, we can return to a prominent figure from the previous chapter, namely Hubert Dreyfus. Dreyfus understands the knowledge used in skill to be procedural (indeed, he seems to equate this directly with KH), understanding this to amount to the experience and intuition gained through repeated practice and exposure that underlies good performance. Across his corpus, Dreyfus puts forward a view of skillful coping that is non-propositional, non-rational, non-conceptual, non-representational and non-linguistic (see H. L. Dreyfus, 2002, 2005, 2007; Dreyfus &

Kelly, 2007). In this way, Dreyfus' view is more 'extreme' than Ryle's view (though he sees their work as being aligned), since his view relies primarily on what he calls 'skillful coping' and appropriate responsiveness more than any manner of 'intelligent' oversight of one's action.

Dreyfus understands himself to be combatting a general intellectualist trend (that extended into the AI and information-processing domains) of thinking that the human mind functions essentially like a computer – it gathers pieces of information and turns them into symbols or representation, relates them to each other, and uses logic, rules, or programs to infer further facts, make decisions, etc. On this line of thinking, an expert performing skillfully must be applying rules and calculations in their mind, whether consciously or unconsciously, in order to perform the way they do. However, Dreyfus argues that “the claim that the expert is using compiled rules is like claiming that because as children we once needed training wheels to ride bicycles, as accomplished bicyclists we must be using *invisible* training wheels” (H. L. Dreyfus, 1997, p. 25). Rather than making the training wheels 'invisible', Dreyfus suggests the expert simply discards them. In the same way that, when prompted, one might be able to go up to the attic and produce the training wheels one used to learn, one could also recall rules that one had learned for how to ride a bike and use them as part of an explanation of how or why one acted as one did. But this does not mean that one was using those rules while performing.

Following his interpretations of Merleau-Ponty, Dreyfus suggests that the relation between knowledge and skill is one where skill doesn't require any reflection or active comparison with memory (i.e. stored knowledge) (H. L. Dreyfus, 2002). Though he does

believe that past experience influences present and future behavior, it is not because one is drawing up stored memories and using them in the process of calculating what needs to be done. Dreyfus understands this latter line of thinking to be an empiricist tendency which suggests that being in a situation similar to a previous one would call forth those past memories to help ‘fill out’ the current situation and help the performer engage. But Dreyfus (borrowing a critique made by Merleau-Ponty) points out that this would create a ‘problem of association’: there are many dimensions along which experiences can be similar to each other and it’s not clear why a given memory would be called up as the appropriate one. The empiricist must then think that experiences are associated with each other or are connected to each other through some act of understanding. In other words, the association/connection between them must be one actively drawn by the subject. This is how the empiricist ends up believing that a subject must draw up some representation of experiences and make connections between them.

However, in line with his previous explications of the stages of skill acquisition, the picture Dreyfus suggests (drawn from his reading of Merleau-Ponty) is one where no representations are needed. Instead,

“[w]hat one has learned appears in the way the world shows up; it is not represented in the mind and added on to the present experience. That is, according to Merleau-Ponty, what the learner acquires through experience is not *represented* in the mind at all but is *presented* to the learner as a more and more finely discriminated situation, which then solicits a more and more refined response” (ibid, 373, original emphasis).

In other words, rather than have skill represented as past experiences and including them in calculations about how to perform given current factors, Dreyfus suggests a picture whereby skill (and the past experience that leads to it) changes the way

the practitioner (literally) sees the world and what possibilities they see as available to them. An expert would be able to see options and possibilities that a novice would not because they are now in tune with parts of their world they weren't before. They act on the 'solicitations' of 'affordances'³¹ they see intuitively – they do not need to discern and identify which possibilities are available or decide which one to take. Indeed, Dreyfus says that “responding to affordances does not require noticing them. Indeed, to best respond to affordances (whether animal or social, prelinguistic or linguistic) one must not notice them *as* affordances, but, rather, as Heidegger says, they ‘withdraw’ and we simply ‘press into’ them” (Dreyfus, 2005, 56).

Take for example, reaching for a doorknob in order to leave the room: in skillful coping, one doesn't need to attend to the doorknob, attend to it *as* a doorknob, or notice the affordance of its turning to open the latch and thus the door. Instead, the affordance is ‘transparent’ and ‘withdraws’ into the object: it is clear to the skilled practitioner on the basis of past experience and becomes part of the fluid action of opening the door (Dreyfus, 2007, 361). Rather than passively taking in input and processing it, skill allows the expert to project their experience out into the world such that “the best representation of the world is thus the world itself” (Dreyfus, 2002, 373).

As additional support for this position, Dreyfus points to contemporaneous models of learning called feed forward simulated neural networks. For these systems,

³¹ Dreyfus says that what we have direct access to in the world is not facts, including facts about affordances (for example, the affordance of an apple as edible understood as a ‘fact’ about the apple). He says “what we are directly open to is not rational or even conceptual(...)” (H. L. Dreyfus, 2007, 356) Instead, what we have direct access to in the world is “the affordance’s *solicitations* – such as the attraction of an apple when I’m hungry” (ibid, 357).

experience is not ‘stored’ as memory and then associated with other experiences. Instead, “if given any input, the connections between ‘neurons’ are modified by a trainer so that that input is paired with what the trainer holds to be the appropriate output. Thereafter, similar inputs will produce the same or similar output” (ibid, 374). This model then provides a way of demonstrating how experience “modifies the connection strength between the simulated neurons” in such a way that “new input can then produce output based on past experience without the net having to, or even being able to, retrieve any specific memories” (ibid).

Ultimately, Dreyfus puts forward the following positive account. He contends that skillful coping in the world, which makes up the majority of human behavior, is intentional without needing representations. This is underwritten by the embodied subject’s tendency towards establishing a ‘maximal grip’³² on their environment. Dreyfus uses a few examples to demonstrate this notion, the first being the tendency to grab an object in such a way as to get the best grip on it. Similarly, when looking at something, say a piece of art, one tends to find the best distance from which to take in both the object as a whole as well as its parts (Dreyfus, 2002, 379). However, Dreyfus expands on this idea and explains that error or deviation from maximal grip is recognized by the subject, and they will align themselves in ways so as to get closer to said maximal grip. He provides the example of a plane going off-course and getting feedback from an airport radio beacon.

³² Though Dreyfus attributes this idea to Merleau-Ponty in the *Phenomenology*, I cannot find any use of this phrase in several translations of said text, and so believe it to be Dreyfus’s own distillation of the text.

Let us suppose the plane gets a signal whose intensity corresponds to how far it is off course, and the intensity diminishes as the plane approaches getting back on course. Thus there is no experience of being on the beam. Rather, when the pilot is on the beam there is no experience at all, but the silence that accompanies being on course doesn't mean the beacon isn't continuing to guide the plane. Likewise, in the case of perception, the absence of tension doesn't mean the body isn't being constantly guided by the solicitations. On the contrary, it means that, given past experience in this familiar domain, everything is going exactly the way it should (Dreyfus, 2007, p. 358).

When doing this, the subject is not representing to themselves the goal of getting 'maximal grip' on the object – one is rather acting in a fluid, skillful way in response to the solicitations of the object and environment. The subject will notice when their "situation deviates from some optimal body-environment relationship" and so take steps to alleviate the 'tension of deviation' by getting closer to the optimum position (Dreyfus, 2002, 378). In this way, skillful coping can be "*purposive* without the agent entertaining a *purpose*" (ibid, 379).

Again, Dreyfus uses neural network theories to help substantiate his claims, specifically Walter Freeman's attractor theory. As before, this model stipulates that learning happens as the strength of connections between neurons changes on the basis of experience. However, Freeman's view adds to the story that learning, i.e. adjusting the connections strengths between neurons, requires adjustments that direct activity toward 'attractors' which have been associated with the object that caused the relevant input. 'Attractors' are points on an 'energy landscape', or points that corresponds to "the amount of energy it takes the whole configuration to be in the state" (ibid, 382). This energy landscape is filled with peaks and valleys – points requiring high energy to achieve, and those requiring low energy. Just as it requires a great deal of energy to get

up a hill and much less to get down it, peaks require high energy to achieve, and valleys require low energy. When an activity is undertaken it generates a burst of energy and it will tend towards a valley or a state of low or minimum energy – these are called ‘basins of attraction’.

The aim of learning is to put the system into a specific energy landscape and hone neural connections so that bursts of activity get directed to specific attractors. Once that energy landscape is achieved, “movements are caused that tend to move the brain state closer to the bottom of the nearest basin of attraction” (ibid). We can understand these minimum energy states, or basins of attraction, as corresponding to Dreyfus’ notion of maximal grip – i.e. states where the subject is pulled to minimize tension between themselves and the environment and achieve equilibrium with it. The energy landscape provides the subject with cues as to what would bring them closer to the relevant basins, without telling them where the basin is.

In sum, Dreyfus elaborates a view in which our fundamental level of engagement in the world is one mediated by our bodies and their fluid push and pull between solicitations and maximal grip on the soliciting objects. As skill progresses, reliance on declarative knowledge continuously diminishes. In time, practitioners transition to the use of procedural knowledge which is based in their experience and acquired intuition around a task or environment. This kind of knowledge does not, on Dreyfus’ account, require any active guidance or monitoring, calculation or deliberation, etc. Instead, this procedural knowledge is guided by the subject’s engagement with the world around them: their environment ‘tells’ them when and how to use their skill and their acquired

‘intuition’ helps them see these ‘signs’. As such, Dreyfus puts forward a view where procedural knowledge underlies skill without any reliance on declarative knowledge. But, as we saw in the previous section and will see again momentarily, there are those who feel that accounts like Dreyfus’s leave out important elements of skill and directly reject his suggestion that declarative knowledge is unnecessary for skillful performance.

2.2.2 – Intellectualist account of Skill and knowledge

2.2.2.1 – Stanley and Krakauer

Intellectualist understandings of the relationship between skill and knowledge tend to underscore the necessity of knowledge for the performance of skill – one could not act skillfully without bringing declarative knowledge to bear on one’s performance. One prominent example of this kind of view is that of (Stanley & Krakauer, 2013). They offer an account which doesn’t necessarily put knowledge prior to skill but does suggest that declarative (or factual) knowledge is essential to skill. They make two fundamental claims: the first is that having a skill requires “being trained past a baseline”, and the second is that, for many skills, having it requires knowing factual knowledge about the skill, including knowing what to do to initiate it. I will focus primarily on the latter claim and only discuss the former insofar as it helps support or clarify the latter.

Stanley and Krakauer (henceforth S&K) begin by suggesting that “skilled action is action guided by ongoing accrual and improving application of knowledge of facts about an activity, though skill is not exhausted by such knowledge” (ibid, 2). More specifically, the kind of knowledge necessary for skill is the kind displayed in knowing

what to do to initiate manifestations of that skill³³. This would be factual (or propositional) knowledge to the effect of “activities $x_1 \dots x_n$ could initiate that action” and would apply to a wide variety of actions within the skillset (ibid, 4). S&K suggest that this ultimately amounts to knowing what to do to begin an intentional action and having a generic sort of knowledge along the lines of “for situation, s , x knows what to do to initiate an action of ϕ -ing in s ”³⁴ (ibid, 5).

To help substantiate this idea, S&K reexamine the case of HM and using it to disrupt the usual dichotomy between declarative and procedural knowledge. Recall that HM was able to improve his skill at a motor task though he did not remember the instances of practice that underlay this improvement. However, S&K point out that the traditional interpretation of this case (suggesting the distinction between declarative and procedural knowledge) does not take into account the fact that, every day HM performed the task, he was given “*explicit verbal instruction*, and was able to use that knowledge each time” (original emphasis, Stanley and Krakauer, 2013, 8). That is, each time HM performed the task he was provided with declarative knowledge about what to do and then used it to perform the task: S&K argue that this instruction was crucial to his being able to do so.

³³ They take this to be a key example of the kind of factual knowledge used in skill and suggest that any other factual information used throughout performance (not just in initiating it) would follow the same kinds of rules and patterns. Though they focus primarily on the case of knowing how to initiate an action, they do not mean to limit all factual knowledge in skillful performance to this knowledge.

³⁴ S&K are quick to point out that knowledge of these starting conditions can be very complex. Furthermore, “an expert will have a larger repertoire of starting actions and will know better how to apply them” (ibid).

To back this up, S&K point to follow-up studies, specifically (Roy & Park, 2010), in which patients similar to HM improved their motor performance of novel tool use over several days. Crucially, they only improved when they received explicit instructions from the psychologists about the tool and how to use it *every day*. S&K are quick to point out that “surely, someone who has skill at an activity does not require such instruction” and that ‘healthy’ skilled people must have declarative knowledge deeply built into their skill, for example, how to use a necessary tool.

As such, S&K interpret Roy and Park’s results as a demonstration that what improves in these amnesic patients is mere motor acuity, not their skill per se. Their abilities are adapting but not in ways that manifest their intentionality and so is not skill. S&K conclude that, for someone to be skilled, they must have propositional knowledge combined with motor acuity and so skill cannot be understood as mere motor or physical improvement on its own. In short, motor skill requires both knowledge *and* a non-knowledge-based motor component – only together do they amount to skill.

However, it is worth noting that the very authors S&K cite to prove their point, Roy and Park, clarify in later works that they do not believe their work should be interpreted as suggesting that the kind of skill at play here, namely tool use, is primarily a question of technical reasoning (Roy & Park, 2018). Though they do note that the study S&K cite highlights the significance of declarative knowledge, especially for tool use and situations that the patient is unfamiliar with, it is also compatible with other research suggesting that this knowledge is mediated by procedural knowledge. Indeed, they suggest that their larger corpus highlights the respective roles of both memory systems

for tool use and their interdependence. As such, it seems that S&K have highlighted an aspect of Roy and Park's work but did so without appropriate contextualization.

2.1.2.2 – Stanley and Williamson

Similarly, in a later work (Stanley & Williamson, 2017) lay out their argument about the relation between skill and knowledge. They initially state that skill is connected to an un-reflective kind of knowledge and later clarify as follows: “Our claim about skills is straightforward: it is that *skills are a kind of disposition to know*³⁵. More specifically, to be skilled at the action type of ϕ -ing is to be disposed to form knowledge appropriate for guiding tokens of ϕ -ing” (original emphasis, *ibid*, 715). So, for example, to be skilled at riding a bike is to be disposed to have knowledge appropriate for guiding the action of riding a bike. More specifically, they identify the kind of knowledge at play here as what is usually thought of as KH. However, they argue that KH falls into a larger category of knowing-wh, which includes knowing-when, -why, -whether, etc. Similarly, following their earlier work, they understand all knowledge-wh, including KH, to be propositionally expressed and so subsumed under KT. So, for S&W, skill is ultimately a disposition to know propositions of the right sort that can guide action. However, rather than putting knowledge or skill as prior to the other, S&W say that “skill at ϕ -ing is a state whose

³⁵ There is an ambiguity in S&K's position here that is worth noting. Their view could be implicitly suggesting a distinction between skill and know-how – that is, if skills, for them, are a disposition to know, then skill and know-how will be different (since one is a disposition and one is knowledge). If this is the case, S&W are perhaps unclear as which of the following is true for skill: skills may be a kind of know how that also involve dispositions to know more specific things OR skills may not be a kind of know-how but rather part of some broader category that involves dispositions to know how to do things. As far as I can tell, it is unclear which of these they think is the case.

nature is constituted through the knowledge relation” and so put the question of priority aside (ibid, 721).

Moving on, S&W explain that they understand skills-as-dispositions to be manifested both as knowledge states and as actions. In terms of knowledge states, an elite athlete’s skill involves knowledge about what the trajectory of a ball will be given the arm and hand movements of the person who threw it, and that this is a manifestation of a disposition to know. Similarly, in terms of actions, skill can be manifested as a “rapid and accurate serve in tennis”, and any other such action. In this way, both declarative and procedural knowledge can come into play in skill, but both should be understood as dispositional factual knowledge that guide behavior. S&W do point out that there is more to skill than dispositional knowledge, including perceptual and motor ability (here, perhaps, following the earlier work done by S&K described above). Even so, S&W understand their account of skill as one “entailing the generic claim that skill in ϕ -ing is knowing at the time of action appropriate facts to guiding ϕ -ing” and so place propositional knowledge as the key explanatory feature at play in both manifestations described above (ibid, 717).

Considerations like these lead (Stanley, 2011a) to provide an alternative process of acquisition to Dreyfus’. Specifically, Stanley suggests that, as a novice, one learns various rules about how to act in the given circumstance. With time, however, “practice allows us to move from the initial situation in which we repeatedly have to consult these rules, to skilled action, where we can act directly upon them” (ibid, 183). In other words, both the novice and the expert acquire propositional knowledge and/or learn rules of how

to perform, but the novice must repeatedly remind themselves of, or consult, the rules which inhibits their fluidity. The novice has not yet developed the automatic processes that bring their behavior in line with their knowledge.

An expert, on the other hand, need not remind themselves of, or consult, the propositional knowledge/rules they possess – they “just *implemen[t]* that knowledge in [their] actions. Practice has allowed the automatic mechanisms that are responsible for executing epistemic states (whether dispositional or not) to take over” (ibid, 184). In other words, the expert doesn’t need to consult their knowledge to apply it to their skill – they can, instead, apply it directly to the situation at hand. In a move similar to the one S&K make in separating out motor acuity from knowledge, Stanley draws an important distinction between propositional knowledge and the ‘automatic mechanisms’ that implement it in action. Both the novice and the expert have knowledge, but only the expert has the automatic mechanisms in place to implement said knowledge in behavior. Skill, then, involves both these components.

2.2.2.3: Takeaway from Skill and Knowledge Debate:

In many ways, this debate mirrors many of the previous issues we discussed before between automaticity vs cognitive control and KH vs KT. Both sides are trying to describe relationship between knowledge and the guidance and manifestation of skill. Non-intellectualists like Dreyfus push for the necessity of procedural knowledge based in past experience and action being guided by trying to get to the ‘appropriate’ equilibrium with their environment; one where practitioners have ‘maximal grip’ and can interact at the most levels. Intellectualists like Stanley, Krakauer and Williamson, on the other hand,

push for the primacy of declarative knowledge, especially propositional knowledge, in both the manifestation and guidance of skillful action. For them, one needs to know facts about how to start an action and how to guide it throughout to be understood as having a skill.

Though Dreyfus's view does seem plausible as an explanation of how procedural knowledge can be garnered and utilized to guide performance, his view still falls prey to the sorts of criticism leveled at him in the previous chapter – namely, that his account is too passive and offers little explanation of the role an agent has in guiding their actions. Indeed, as I read Merleau-Ponty (MP), Dreyfus leaves out much of how his account emphasizes the subject's active role in perception and action that happens in conjunction with the more 'passive' elements Dreyfus highlights. Indeed, MP's view emphasizes the subject's active role in shaping their environment in line with their goals and intentions as a key component of the nature of perception in ways that Dreyfus seems to leave out altogether (Merleau-Ponty, 2012). Though this does not necessarily mean that Dreyfus is wrong in his understanding of how procedural knowledge is gained and utilized in performance, it does suggest that his descriptions lack a sense of the kind of guidance or oversight one would have over their own actions and thus leave out (even the potential for) the role of declarative information in assisting in action. As such, his views stray from the sort of hybridity we have been highlighting up till now.

Similarly, Stanley, Krakauer and Williamson's positions seem, on the surface, to be pretty standard intellectualist readings of the relation between skill and knowledge. After all, they all still seem to want to reduce knowledge to facts and propositions,

without which one could not manifest or guide one's skill. Furthermore, we saw that their use of a key study to make their point relies on what the authors interpret to be a misreading of their view – one that suggests that factual or declarative knowledge is primary for skill without acknowledging the role and contributions of procedural knowledge. As such, we might worry that their position is already too far skewed to the 'intellectualist' side of things and thus suffers the same kind of critique we just made for Dreyfus – that is, they may be right in their explanation and highlighting of the role of declarative knowledge and skill, but they do so in a manner that completely negates the contributions of procedural knowledge.

However, their views (especially S&W's) seem to make some concessions to the non-intellectualist camp. For one thing, they highlight the role of dispositions in skill. In terms of our previous discussion of automaticity and the speed at which skill is performed, utilizing the notion of dispositions in skill can be understood as a way of a taking on or conceding to this apparent non-intellectualist strength – if the nature of skill involves dispositions, they don't require careful (and time-consuming) consideration or calculation, but can instead act on a quick and intuitive propensity. In this way, S&K&W's revised positions take on some non-intellectualist intuitions and build them into their view.

Another apparent concession is that S&W allow both procedural and declarative knowledge will be at play in skillful action. In this respect, their view seems in line with the hybrid positions already discussed. However, they continue to maintain that both of

these forms of knowledge can ultimately be understood as different kinds of propositions and so maintain a core intellectualist component. Even so, insofar as their views aim to muddle the distinction between declarative and procedural knowledge and highlight some more basic, shared third feature, their position might still align with the general positions taken up till now.

Borrowing from S&W, we can maintain the basic framework of saying the skill requires the appropriate relation to knowledge. However, let us specify this differently than they did and add in some of our newer criterion from this chapter...

Criterion 5*: skill must have an appropriate relation to knowledge, which includes...

- ...an interdependence between KH and KT. More specifically, KH utilizes elements of control and KT in performance, while KT relies on at least some KH to be able to use, interpret or execute it.
- ... an interdependence between declarative and procedural knowledge. More specifically, procedural knowledge utilizes elements of declarative knowledge in performance, while declarative relies on at least some procedural knowledge in order to be used, interpreted or executed.

2.3 – Virtue Epistemology

Up till now we've focused on what kind of knowledge is at play in skillful performance and pushed for a hybrid understanding between KH and KT, procedural and declarative. However, though we have touched on the borders of this issue, we have not yet focused on another important question: namely, whether knowledge requires skill or

vice versa. That is, though we have discussed whether skill utilizes different kinds of knowledge, we have not yet discussed whether having knowledge relies on having a certain kind of cognitive skill. Getting a grip on this will help take us deeper into the relation between skill and knowledge and help clarify the degree to which skillful action and behavior are basic to our way of being. We will focus on a few key representative key figures in this debate in order to see how this plays out.

2.3.1: 'Non-intlectualist' accounts in Virtue Epistemology:

Within the virtue epistemology debate, one of the most influential voices is that of Ernest Sosa who argues that knowledge is to be understood *in terms of* skill – or in other words, that we should understand knowledge as a type of *apt performance*. Sosa's AAA structure outlines his means for assessing the performance of an action. It can be assessed for accuracy, i.e. whether it is successful in its aim; for adroitness, i.e. whether it manifests skill on the part of the performer; and lastly for aptness, i.e. whether the success of the action is due to the performer's skill (Sosa, 2007). Our purposes require special attention to the final condition of aptness, which encompasses the other two criteria.

For Sosa, a performance is apt if it is both successful *and* successful because of one's skill. His famous example is that of an archer. Suppose an untrained archer looses their arrow rather spuriously but it is blown back on course by a gust of wind and hits the bullseye. In this case, the archer is successful, but their success has little to do with their skill. Their success is due, instead, to the wind's benevolent influence. This, then, is not a case of 'apt' performance. Similarly, consider a skilled archer who looses their arrow on

target but a gust of wind first blows it off course and then back on course and finally it hits the bullseye. In this case, too, we might think that the archer's success has little to do with their skill despite the fact that we know they are skilled. The influence of the wind seems largely responsible for the success of the shot. This second case, then, is also not one of 'apt' performance.

Sosa suggests that any performance with an aim could be subject to the AAA structure, including those without an intentional aim, such as a heartbeat – “a heart succeeds if it helps pump blood, even absent any intentional aim” (23). On this line of thought, beliefs, too, could be understood as performances and so fall under the AAA structure. In other words, beliefs might also be performances without an intentional aim, but could still be relevantly assessed using the AAA method. Sosa says that we can “distinguish between a belief's accuracy, i.e. it's truth; its adroitness, i.e., its manifesting epistemic virtue or competence; and its aptness, i.e., its being true *because* competent” (23).

From here, Sosa states that knowledge entails belief – but a special kind of belief. In order to explain this, he introduces a few key requirements for belief. The first is the idea of 'safety', which requires “that not easily would [a belief] fail by being false, or untrue. A belief that p is *safe* provided it would have been held only if (mostly likely) p” (25). The second is the notion of 'sensitivity', which requires that “were it not so that p, [a subject] would not (likely) belief that p” (25). Sosa argues that, though it is often mistaken for sensitivity, safety is the most relevant concept for knowledge and should be more carefully and separately identified. As such, he identifies a particular iteration of

safety as his target. This he calls ‘basis-relative safety’. He says that “[w]hat is required of one’s belief, if it is to constitute knowledge, is at most its having some basis that it would not easily have had unless true, some basis that it would (likely) have had only if true” (26). In other words, rather than merely focusing on the plausibility of the belief itself, Sosa focuses on whether *the basis* of the belief is plausible. In this case, a belief would constitute knowledge if it has a solid enough grounding on the basis of which it will be reliably true.

With all this in hand, Sosa then concludes that knowledge based on basis-relative safe beliefs is also subject to the AAA criteria and, specifically, assessment by the aptness criterion. Sosa says that “knowledge is just a special case of (...) creditable, apt performance” (31). Knowledge, then is underwritten by apt, basis-relative safe beliefs, which are part of our epistemic competences. Sosa understands a competence to be “a disposition, one with a basis resident in the competent agent, one that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it” (29). This competence is a form of skill, and so, for Sosa, skill underwrites knowledge³⁶.

Sosa thinks that most of the time, in ordinary conditions, our (ordinary) knowledge is reasonably safe, and this is important for uptake of *deliverances*. Sosa

³⁶ Throughout the book from which this explanation is drawn, Sosa is clear that he is differentiating between at least two types of knowledge – what he calls ‘animal’ knowledge and reflective knowledge. Reflective knowledge is the more demanding of the two, but Sosa thinks both involve apt belief. ‘Animal knowledge’ only requires apt belief in the way that has been described above. Reflective knowledge, on the other hand, requires both apt belief and *defensibly* apt belief (i.e. apt belief that the subject aptly believes to be apt). This latter condition means that the subject can defend the aptness of their belief against skeptical doubt (pg. 24)

understands deliverances as “*any particular delivering of a certain propositional content*”, either from memory, perception, conclusions, intuitions, etc. (original emphasis, 101). Under usual circumstances and/or appropriate conditions, deliverances are safe outright, because the conditions under which the content was gathered are (most likely) true and there is no reason to doubt the appropriateness of the conditions. Unless we have some reason to doubt the circumstances, most deliverances are accepted by the subject, and so safety serves a key role in our most basic intake of information and is a key grounding of knowledge. However, safety does not account for the *correctness* of deliverances or beliefs – rather, this has to do with one’s competences. Sosa says that “[a]cceptance of a deliverance thereby constitutes knowledge only if the source is reliable, and operates in its appropriate conditions, so that the deliverance is safe, while the correctness of one’s acceptance is attributable to one’s epistemic competence” (103).

One reason that Sosa suggests that we should understand knowledge as an instance of skill is because it helps explain what is happening in Gettier cases – i.e., cases in which the subject has a belief that is both justified and true but falls short of knowledge. Consider the following example from Gettier, himself: Suppose that Smith and Jones applied for the same job. Smith was told by the president of the company that Jones would be selected for the job and he also saw Jones counting out the 10 coins in his pocket before going into the interview. Because of these pieces of evidence, Smith has a justified belief that “the man who will get the job has ten coins in his pocket”. However, as it turns out, Smith himself gets the job and Smith also has 10 coins in his pocket that he didn’t realize were there. So, Smith has a justified belief that “the man who will get

the job has ten coins in his pocket”, based on the evidence he had beforehand – and, as it happens, it is also true despite it being about the wrong person. However, we can clearly say that this belief doesn’t amount to knowledge because he seems to have ended up at this correct conclusion accidentally. (Gettier, 1966)

For Sosa, in cases like this one, the subject may have a justified, true belief, but one that doesn’t amount to knowledge because it is not apt: i.e., its success is not due to the competences of the subject. In other words, Smith has a belief that is true (after all, the man with ten coins in his pocket *did* get the job) and justified (on the basis of the evidence he had, including the testimony of the president). As such, Smith formed a belief on the basis of his general epistemic competences. Even so, this belief does not constitute knowledge because Smith’s usual belief forming competences are not responsible for the truth of his belief. Indeed, Sosa says that Smith’s reasoning on the basis of testimony from the company president “does of course help explain why the believer has that belief, but it does not in the slightest help explain its *correctness*” (Sosa, 2007, 96). Like the case of the skillful archer, Smith skillfully deployed his belief-arrow, and it hit the bullseye – but it hit the bullseye for reasons other than his own skill. Sosa suggests, then, that we can understand this, and other Gettier cases, as failures of apt belief, and so as failures of knowledge.

2.3.3: an “Intellectualist” Account of Virtue Epistemology

However, there have been several critiques of this position based on demonstrating that understanding knowledge as apt belief doesn’t account for other forms of Gettier cases. For example, (Miracchi, 2015) suggests that this doesn’t explain

Chisholm's Gettier case (Chisholm, 1989). The case is as follows: Suppose Annette is walking through a field and seems to see a sheep in front of her – so, she believes that there's a sheep in the field. However, it turns out that what she's really looking at is a sheepdog. Even so, there is truly a sheep behind a rock, out of sight, that the dog is looking after. It also happens that, unbeknownst to Annette, sheepdogs in this area almost always accompany sheep in the field. In this case, Annette's belief (that there is a sheep in the field) is true – there is one. Similarly, her belief is justified since her perceptual experience is generally reliable and, let's assume, the visual differences between a sheep and a sheepdog at a distance are minimal. Even so, one wouldn't intuitively consider this a case of knowledge: Annette doesn't *know* that there is a sheep in the field.

However, Miracchi suggests that, on Sosa's account, Annette would have an apt belief and so, also, knowledge. This is because "Annette's competence to believe truly that there is a sheep in the field on the basis of a perceptual experience as of there being a sheep in the field is causally responsible, not just for the existence of her belief, but also for its truth" (34). In other words, it is Annette's competence at true belief formation on the basis of perceptual experience that explains its success. Miracchi seems to have two reasons for thinking this. First, she thinks that Annette is being *systematically* Gettiered in this case: as was already stipulated, sheep are almost always accompanied by sheepdogs in this area and so Annette only sees the sheepdog because the sheep is there, too. As such, her exercise of successful belief-competence (i.e. that there is a sheep in the field based on having seemingly seen a sheep) is causally dependent on the fact that she sees the sheepdog. If she hadn't seen the dog, she wouldn't have formed the belief that it

was a sheep: the true belief depends on her having formed a belief based on (what turns out to be faulty) perceptual experience, but the belief-forming competence process works as it should.

The second reason Miracchi suggests that, on Sosa's account, Annette would have knowledge is that, since Annette's belief forming process is working as it should, she doesn't notice the subtle differences between the sheepdog and a sheep that she might have otherwise noticed if she was being more careful. That is, since she has not identified any reason to doubt her usual conditions or competences, she allows it to run normally and doesn't notice the minor details that would prove her beliefs unsuccessful. If she had, she would have noticed that what she was looking at was really a sheepdog and, as a result, would have formed the incorrect belief that there were no sheep in the field. So, Miracchi says that Annette's "competence thus *is* 'a factor that, either singly or in combination with other factors, accounts for how the belief is true rather than false' (Sosa, 2007, p 96), and so her believing truly is causally explained by her competence" (Miracchi, 2015, 34).

Miracchi considers some defenses available to Sosa to counteract this conclusion. The first has to do with the dispositional nature of competences and their *characteristic manifestations* (Sosa, 2010). Under normal circumstances, Sosa says that a subject will have a true belief as long as the characteristic manifestation of their disposition to believe truly (i.e. their competence) is actually manifested. However, this doesn't happen in Gettier cases – that is, whatever the characteristic manifestations of the disposition are, they are not actually achieved in Gettier cases. Furthermore, in (Sosa, 2015) Sosa

explains that what's happening in Gettier cases is not a characteristic manifestation of their disposition to believe truly, but a *mimicking* of said disposition. That is, in Gettier cases, a subject does believe justifiably and truly and their disposition to believe truly (i.e. their competence) is causally involved in this being the case, but the relation between the competence and the belief is not the right one: it isn't the characteristic manifestation which, in this case, is only being mimicked.

Even so, Miracchi suggests that these responses fail because they still do not account for the case with Annette described above. The systematicity of the Gettier-ing in this case makes it so that this *is* an instance of the characteristic manifestation of Annette's competences. As Miracchi puts it, "there is no basis for denying that Annette manifests a disposition to believe truths on the basis of having a perceptual experience as of a sheep in the field, although she fails to manifest other dispositions normally associated with such subjects, such as a disposition to believe truths on the basis of seeing things as they are" (37).

In light of this, Miracchi seems to lean towards intellectualist positions like S&W's and suggest a knowledge-first account. She explains that, on a view like Sosa's, the competences/dispositions that explain knowledge are *competences to believe truly*. Beliefs are the exercises of this competence which may or may not be true (i.e. successful). However, on Miracchi's view, the relevant competences/dispositions for knowledge are *competences to know*. Instances of knowledge are the exercises of this competence. That is, rather than Sosa's picture where a probabilistic relation between belief and truth determines knowledge, Miracchi offers a picture where there is a

probabilistic relationship between “the subpersonal cognitive and environmental facts” that determines knowledge (ibid, 53). In terms of the case with Annette, Miracchi’s view suggests that Annette doesn’t in fact know that there is a sheep in the field (which is the intuitive answer), because the relationship between the facts of her subpersonal cognition and her environment do not align and so this is not an instance of a manifestation of her competence to know.

However, this means that, competence has a circular relation with knowledge – one where “what it is to be knowledge is to be a manifestation of a competence to know, and what it is to be a competence to know is to be a competence whose manifestations are particular cases of knowledge” (ibid). But Miracchi doesn’t think this circularity is vicious – instead it points towards a co-dependence between competence and knowledge. As such, Miracchi’s view is a more nuanced intellectualist view that suggests turning the conversation away from beliefs and towards knowledge, itself, as central, but also seems to reject the question of priority between skill and knowledge.

2.3.2: Takeaways from Virtue Epistemology Debate:

There are a few things I think it worthwhile to take away from Sosa’s analysis. For one, his has been the first so far to examine success conditions for skill and ruling out chance or accident. As such, I suggest the following criterion for our list:

Criterion 6: skill should be understood to have success conditions which include achieving successful performance *due to one’s own skill* (i.e. aptness) as opposed to any kind of accident or unintentional outcome.

However, as far as our discussion goes regarding whether skill is prior to knowledge or vice versa, more needs to be said. Sosa's account rightly highlights aptness (i.e. success due to one's own skill in both knowledge formation and skill more generally). Miracchi seems to take this on, too, in so far as her view maintains the role of competence (though she highlights a different kind of competence than Sosa). Similarly, Sosa emphasizes the notion of the safety of knowledge formation as it is tied to environmental factors. Miracchi seems to explicitly take this on, as well, though she highlights subpersonal cognitive processes in ways that Sosa does not. So far, then, there's a great deal of agreement between the two parties.

However, our commitments from previous analyses should push us closer towards a position that highlights an interrelation or a codependence between the relevant notions at play. In that regard, taking on a view like Miracchi's, which does exactly this, will be in line with much of what has been already established. Her view is 'intellectualist' insofar as it highlights a *competence to know* (as opposed to a *competence to believe truly*) as central to knowledge, but, as highlighted above, takes on many aspects of the non-intellectualist position. Indeed, her view merely redirects the conversation away from belief and towards knowledge: it doesn't undermine the role of skill. The 'circularity' of her view further demonstrates its hybridity.

As such, I suggest this merits another addendum to an earlier criterion about the relation between skill and knowledge, namely:

Criterion 5:** skill must have an appropriate relation to knowledge, which includes...

- ...an interdependence between KH and KT. More specifically, KH utilizes elements of control and KT in performance, while KT relies on at least some KH to be able to use, interpret or execute it.
- ... an interdependence between declarative and procedural knowledge. More specifically, procedural knowledge utilizes elements of declarative knowledge in performance, while declarative relies on at least some procedural knowledge in order to be used, interpreted or executed.
- ... and interdependence between skill and knowledge – each utilizes and relies on the other.

2.3.3: Interlude: Safety, Flexibility and RRR:

However, there is another notion that we can take from Sosa's analysis which will also influence our previous conversation around flexibility: namely 'safety'. Recall that, in the previous chapter, we understood flexibility as either the ability to adapt the performance of a skill across various circumstances and still perform well or being able to achieve the same goal through a variety of uses of skill. In the first case, the example used was of playing basketball both indoors and out or playing with a heavier/lighter ball than usual, etc. In the second case, the example used was that of a chef cooking a meal – the same dish could be created through any number of iterations of the same steps. A skill that could take into account the various challenges and differing influences at play in these respective scenarios would be a flexible one. Recall, also, that a safe skill was

understood to be one which was not likely to fail under normal circumstances. There is some overlap between these two notions (flexibility and safety), but also ways in which they can be teased apart. As such, I'd like to spend some time examining them so we can get a better understanding of our second criterion (about F&F).

To begin with, let's examine this apparent overlap. There might be an initial sense in which these two terms denote a similar notion, i.e. that a skill be successful across a wide variety of circumstances. However, I suggest that the two terms are actually trying to highlight different things and would have different metrics of performance. On the one hand, what the notion of flexibility seems to be underlining is that a flexible skill is one that is *resilient*³⁷ across circumstances, i.e. one that is adaptable across various iterations, challenges, scenarios, etc. In general, the notions of flexibility and resilience are latching onto categories of performance such as ideal/normal performance, non-ideal/subpar performance, and failure. Safety, on the other hand, states that performance should not be likely to fail under normal circumstances. In other words, what the notion of safety seems to be getting after is that the skill is *reliable* within its own domain. The metrics for performance in this case have more to do with normal performance, performance within a given range, and performance outside the domain of the skill in question³⁸. In this way, the two concepts come apart and help highlight an important conceptual difference for skill.

³⁷ This, and the other terms that follow, are ones that I have chosen and that are used slightly differently in different literatures. However, I mean to use them in a sense that is devoid of other associations and tied to what I think are their most intuitive meanings.

³⁸ If that skill is able to be successfully performed or used outside of its intended or initial domain, we could then say that it was *transferable*.

When one's skill is both resilient and reliable, i.e. both flexible and safe, we might consider it a *robust* skill, i.e. one that maintains a certain internal integrity across circumstances. That is, as a result of being both resilient and reliable, we would consider it robust. For example, suppose basketball player A can play well both indoors and outdoors, i.e. their skill is both a) resilient to the change in environment and b) reliable in both these circumstances. They have a more robust skill than player B who only plays well indoors, i.e. whose skill is a) not resilient and b) is only reliable in one circumstance. Similarly, both players A and B would have a more robust set of skills still than player C whose skill is neither resilient nor reliable in either circumstance.

Given this analysis (and inspired by Sosa's AAA account described above), I propose an addendum to our second criterion regarding fluency and flexibility. The notions of F&F can remain largely the same as they were before, i.e. that F&F will utilize both automatic and cognitive control. However, using an RRR account, we now have a means of assessing and describing F&F that we didn't before. Even so, unlike Sosa, whose AAA account seems most relevant to individual acts of performance, this RRR account applies most directly to act types or general performance over time. Robustness, then, unlike adroitness, has built into the notion that is only relevant across a range of circumstances and performances, rather than individual instances.

However, there may be reason to expand this system a bit further. Suppose a juggler is able to successfully perform the same kind of juggle (say, juggling in a circle) in many different circumstances: indoors, outdoors on a windy day, standing in the pool, etc. Given our thinking above, we say that his skill is resilient and reliable across many

situations. However, suppose our juggler has several different ways of juggling (juggling in circles, in crosses, at varying speeds, etc.) and which one they perform depends on the circumstance. For example, suppose they'd usually default to juggling in circles but, given the influence of the wind, they juggle in crosses, and standing in a pool they vary their speed to match the 'waves' hitting them, etc. In all these different circumstances, their skill is successful and could be understood as resilient, reliable, and robust. However, the latter case seems to suggest a different notion than was originally highlighted, namely responsiveness to circumstance. That is, the juggler's skill includes a repertoire of throwing patterns which can be utilized depending on their appropriateness to the circumstance, which suggests an adaptability to their skill not previously captured by resilience and reliability. So, perhaps this is really an RRRR account: where resilience, reliability and responsiveness help determine the robustness of the skill.

With all of this in mind, I propose an addendum to our second criterion as follows:

Criterion 2*: skill is typically manifested in fluent and flexible action.

- The mechanisms underpinning fluency and flexibility will utilize both automaticity and cognitive control throughout performance in ways that are perhaps ambiguous and difficult to separate.
- The fluency and flexibility of skill can be assessed using the RRRR method: i.e. assessing its resilience, i.e. adaptability across various iterations, challenges, scenarios, etc; reliability; i.e. maintenance of consistency across circumstance; responsiveness, i.e. adaptiveness to

changes in circumstance; robustness, i.e. a skill that is resilient, reliable and responsive.

2.5 Conclusion:

This brings us to the end of our investigation of the relevant epistemological literature. Having gone through it all, we have gleaned our full list of criteria which I collect below. With this in hand, we will now move onto an investigation of a case study to help illustrate the kinds of skill that will be of particular interest as we move on to later chapters, namely walking.

Criteria of Skill:

Criterion 1: Skill is learned, acquired or gained through practice. It is not innate ability

Criterion 2: skill is typically manifested in fluent and flexible action.

- The mechanisms underpinning fluency and flexibility will utilize both automaticity and cognitive control throughout performance in ways that are perhaps ambiguous and difficult to separate.
- The fluency and flexibility of skill can be assessed using the RRRR method: i.e. assessing its resilience, i.e. adaptability and transferability across various iterations, challenges, scenarios, etc; reliability; i.e. maintenance of consistency across circumstance; responsiveness, i.e. adaptiveness to changes in circumstance; robustness, i.e. a skill that is resilient, reliable and responsive.

Criterion 3: skill is aimed at ‘getting it right’ and is underscored by the practitioner’s intentionality.

Criterion 4: to be understood as skilled entails performance above a context specific threshold.

Criterion 5:** skill must have an appropriate relation to knowledge, which includes...

- ...an interdependence between KH and KT. More specifically, KH utilizes elements of control and KT in performance, while KT relies on at least some KH to be able to use, interpret or execute it.

- ... an interdependence between declarative and procedural knowledge.

More specifically, procedural knowledge utilizes elements of declarative knowledge in performance, while declarative relies on at least some procedural knowledge in order to be used, interpreted or executed.

- ... and interdependence between skill and knowledge – each utilizes and relies on the other.

Criterion 6: skill should be understood to have success conditions which include achieving successful performance *due to one's own skill* (i.e. aptness) as opposed to any kind of accident or unintentional outcome.

Chapter 3

Walking as Skill:

3.0: Introduction:

When asked to think of physical skills, researchers (especially in philosophy) tend to reach for a few key paradigms of skill, such as playing an instrument or playing a sport. It is easy to see how these kinds of skills fit the mold of everything discussed in the previous chapters – indeed all our gathered criteria for skill were drawn from examples of this sort. However, I want to incorporate much more basic skills into our conversation, as well – skills like walking. By doing this, we can focus on the most basic level of engagement with an environment and how these engagements set the foundation for all the other skills that rely on them. Indeed, these basic skills probably make up the vast majority of our skills and are the most crucial for us being able to act normally on a day-to-day basis. Getting clear on this kind of action as a skill helps highlight the core elements of these skills and helps us avoid focusing on potential examples that might widen the scope of investigation too far (keeping in mind our aim of focusing on physical skill).

Once we get to this kind of activity, however, there might be worries about whether they are actually skills or innate abilities. For example, while discussing complex skill, S&K (2013) note the fact that complex skills will be made up of more basic skills. They use the example of someone who is skilled at tennis having a more basic skill of serving. However, they point out that, at some point, we would bottom out at basic actions, such as grasping an object or lifting one's arm, which should not be understood as skill but rather basic action. This is because these activities “are not acquired by or

improved upon by training in adult life” (Stanley & Krakauer, 2013, p. 5). For S&K, something like walking might be understood to be a basic action for the same reasons. Similarly, we might not usually think of something as simple as walking as a skill because it seems something so basic, so natural, something that the vast majority of humans do – or, as S&K suggest, it isn’t acquired or improved upon with training during our adult lives.

In what follows I argue that walking meets all the criteria of skill discussed in the previous chapters and follows the same general patterns. As such, I mean to demonstrate, first, that some even very basic activities should be understood as skills and, second, to question the framework of adult practice as necessary for an action to be a skill. Lastly, I will use walking and other examples of this sort, i.e. of very basic skillful activity, in the succeeding chapters and their relevance to procedural memory. As such, getting clear on their nature in detail early on will be critical.

To do this, I examine a comprehensive review article by Adolph and Robinson (2013) (henceforth referred to as A&R) about the acquisition of walking in infants. A&R’s aim in their review is slightly orthogonal but complimentary to our purposes here. They aim to dispel a central metaphor about how walking develops in infants, namely the ‘milestone’ metaphor. This picture is one in which learning to walk is a standard developmental ‘milestone’ along a linear chain of progressive milestones, each a prerequisite for the next, which children must pass through in the ‘correct’ order. Thinking this way demands that there are steps along the chain that cannot be skipped, occur simultaneously, or in different orders. This interpretation lends itself to a view of

walking as ‘innate’ and a simple natural progression through these various stages of development³⁹. However, A&R compiled significant research that demonstrates this not to be the case. In fact, infants demonstrate clear learning processes in the acquisition of walking which rely on repeated attempts at practice. At each stage of learning, infants become increasingly fluent and flexible with their skills and open up further avenues of locomotion. Furthermore, this process is not a linear one - infants often flout all the ordinances described in the ‘milestone’ metaphor above: they skip over some phases, learn various methods of locomotion simultaneously, and learn them in the ‘wrong’ order.

In addition to helping ground the notion that even basic actions like walking can/should be understood as skill, this literature also serves to add a new dimension to our conversation about the acquisition of skill. In previous chapters, we focused on the extent to which automaticity or cognitive control was at play in reaching higher levels of performance, as well as well as what kind of (relationship with) knowledge is at play during skillful performance. However, this literature can help focus us on how instances of practice can help lead from one stage of skill to another, and how a web of foundations is laid for further progress. That is, it introduces a conversation about the developmental acquisition of skill which has not yet been discussed.

I will begin by focusing on what I believe to be the most contentious element in determining whether walking is a skill or not, namely whether and how it is *learned*

³⁹ Though it need not – this view could still understand walking as an ability that is learned (and so, a skill), but the process of learning it is a necessarily linear one. This latter view is, in fact, the interpretation I understand to be prominent amongst psychologists with this viewpoint.

(Criterion 1). This demonstration also opens up space to discuss the intentionality of walking (Criterion 3). From there, I will discuss the ways in which walking demonstrates the kind of fluency and flexibility highlighted in previous chapters (Criterion 2), and how this demonstrates performance above a given threshold (Criterion 4), as well as the aptness of the practitioner (Criterion 6). Finally, I will discuss the relationship between walking and knowledge (Criterion 5) and how walking may flesh out our previous conversation about the relation between skill and knowledge and opens up a new framework while maintaining our previously established points.

3.1 – Learning to Walk

As we saw in previous chapters, skill is usually thought to be learned, i.e. gradually acquired over time and practice such that one progressively becomes increasingly proficient with the action/task in question and eventually becomes fluent and flexible with it. Walking can be shown to be acquired in just this way: infants are not born able to walk, and their being able to do so is not simply a question of reaching the right biomechanical arrangements with age and development. Though there are biomechanical contributors that must be in place for an infant to walk, these are not enough: the infant must actually try to move themselves and get around an area, thereby finding and practicing various forms of locomotion. Indeed, infants often fail to move in various ways, and either keep trying or try something else to move them around. As these methods become increasingly competent, they lay the groundwork for further developments, which in turn lay the groundwork for the next bout, and so on until they gradually build their bodily competency to walk. Ultimately, most infants will come to

walking through varied motor methods and trajectories, suggesting that there is no one correct developmental process to walking and that each infant finds their own idiosyncratic way of getting there based on their own personal experience and practice. To establish this, we will explore a great deal of data presented by A&R in their review.

3.1.1 – Starting points for walking

While many understandings of skill in developmental psychology suggest that skills have a distinct starting point, A&R begin by endorsing a view of skill that does not require one. Of course, any observation of skill, conceptual or empirical, will begin somewhere, but they say this is “merely a convenience”; for A&R, every point of development is based in an earlier point (Adolph & Robinson, 2013, 3). To demonstrate this, they consider the following traditional understanding of the main starting point for walking; namely leg alternation (Grillner & Wallen, 2004; Spelke & Newport, 1998). Indeed, a particular manner of leg alternation (50% phasing) is the defining characteristic of walking as a gait pattern and is demonstrated by infants long before they can walk. They demonstrate leg alternation in the womb, lying on their backs, when held up above a horizontal base (like a table), when held in midair, when held up against a vertical wall and even held upside down up to a ceiling (Andre-Thomas & Autgaerden, 1966; de Vries et al., 1982, 1985; Piontelli, 2010; Esther Thelen, 1979; Esther Thelen & Fisher, 1982; Ulrich, 1989). However, A&R point towards two reasons to think that leg alternation is not the starting point for walking. First, they demonstrate that alternating leg movements, themselves, have a developmental history. Before fetuses start moving their legs in alternation, they move their legs in synchrony (Nakayama et al., 2001). Similarly,

moving legs in synchrony is underpinned by even earlier development, namely of feedback from proprioception about the position and movement patterns of limbs. Even before this, fetuses demonstrate spontaneous leg movements that contribute to the development of their nervous systems, muscles and bones (Moessinger, 1983). All of this suggests that even the supposed ‘starting point’ is dependent on earlier ‘starting points’ and so none is likely the ‘core starting point’ for walking.

Secondly, A&R point out that fetuses and neonates spontaneously move their legs in any number of ways besides alternation. For example, they might do “single leg kicks, symmetrical double leg jumps, and asymmetrical alternation that differs from 50% phasing” (A&R, 2013, 5; discussing (Suzuki & Yamamuro, 1985). Indeed, at different points of development, the exhibition of one or more of these other leg movements is often more common than leg alternation. It is unclear, then, why leg alternation is so favored amongst all these various movements, especially when, as A&R point out, mere alternation of the legs is far from enough to be able to walk. Indeed, walking requires being able to speed up or slow down, steer and navigate through/around obstacles, adjusting to changing in footing, slopes, or the lode one may be carrying, and maintenance of balance throughout. All of this will require “flexibility of interlimb coordination, perceptual guidance, postural control, muscle strength, and the motivation to go somewhere” (A&R 2013, 5), each of which have their own developmental history and none of which, alone, would underpin the entire ability of walking.

In general, though, it seems that many resources must be in place before an infant is ready to walk (Esther Thelen, 1995; Esther Thelen & Smith, 1996)). This will include

various patterns of leg movements, such as leg alternation, but also the development of neural systems that underpin “pattern generating circuitry in the spinal cord, dedicated executive regions in the midbrain and motor cortex, and integrative networks necessary for coordinating and fine motor control in the cerebellum” (A&R, 2013, 13). Similarly, though sensitivity to visual flow begins long before mobility, it improves dramatically over the first year of life and is crucial for mobility (B. I. Bertenthal et al., 1997, 2000; Gilmore et al., 2004; Gilmore & Rettke, 2003). Furthermore, bodily dimensions that transition from a top-heavy to a more cylindrical one, as well as developing sufficient muscle strength and balance control, all seem to be necessary contributors to walking (McGraw, 1945; E. Thelen et al., 1984; Woollacott et al., 1989). However, even with all this in place, these are only a basic foundation and are not a substitute for actual, physical practice with self-generated movements.

3.1.2 – Experience:

Though age and experience are often correlated (i.e. older infants tend to be more experienced), experience, itself, is the best predictor of improvements in mobility and can be decoupled from age (Karen E. Adolph, 2000; Karen E. Adolph et al., 2003; Karen E Adolph et al., 1997). However, experience doesn’t seem to operate in a linear way on smaller scales – that is, “new skills stutter in and out of infants’ repertoires from day to day and old skills slowly die out, meaning that experience is not a continuous function of time” (A&R, 2013, 21). Experience will be at least partially determined by the amount of practice an infant has and the amount of exposure they have to particular events or situations. In general, infants get extremely variable and dispersed practice with walking

(Karen E. Adolph et al., 2012). The average 14-month-old toddler takes about 2000 steps an hour, travels 2520ft (~half a mile) and falls about 15 times. They usually walk in short, interspersed bouts with long breaks in between, each usually in a different social and physical context. Even so, quantitative data like this is far from being lined up directly with the necessary and sufficient conditions for experience in skill.

What has been thought to be important for infants' experience and so learning to walk is continuously pushing themselves onwards at the edge of their current skillset. Indeed, infants keep pushing their mobility at every stage: for example, rather than 'getting stuck' at crawling which, given the right amount of experience, works perfectly well to get them around, they push themselves onward to find another means of mobility. This idea fits in with a famous notion within developmental psychology, namely Vygotsky's Zone of Proximal Development (ZPD), which identifies "the range of performance between what children can do currently and the level they are likely to achieve next" (A&R, 2013, 22; discussing (Vygotsky, 1978). Another complimentary concept in developmental psychology is "scaffolding", in which "support from an external source (typically a more knowledgeable person or a cultural tool) can move children through their ZPD to the next level"⁴⁰ (ibid).

Along these lines, then, for infants to move from one mobility benchmark to another, they must push the current limits of their abilities until they slide into the next one. As we shall see, to do this, they utilize whatever resources they can find, including

⁴⁰ It is perhaps interesting to note that this is in line with Csikszentmihalyi's understanding of flow and needing to push to the edges of one's skill set to help get to it, whatever one's skill set may be, as opposed to Dreyfus's account.

the full extent of their space, furniture and willing adults around them. Each of these bits of scaffolding can help propel an infant from one benchmark to another but, ultimately, it is the infants' experience, understood as a sort of aggregate understanding based on past interactions, that helps set the 'limits' and give them the means to push past them. This experience continues to accumulate but, as we shall see, this process is not always as linear as it seems.

3.1.3 – From crawling to cruising to walking:

Before infants can walk, they pass through several means of mobility. Their selection of solutions for movement seems to depend largely on functionality: whatever works, works. Even so, each of these forms contributes to their being able to perform whatever they do next by shoring up competencies and abilities necessary for their performance and helps demonstrate their process of moving through various ZPD's along the way. The kind of trial and error used by infants, as well as their failures which spur new tactics or, in other words, the gradual accumulation of personal experience, strongly suggests that walking is learned rather than inbuilt. To demonstrate this notion, A&R take the reader through a great deal of data on the 'benchmarks' passed on the way to walking.

The iconic pre-walking benchmark is generally depicted to be crawling on hands and knees. However, infants employ many additional forms of crawling including the 'belly crawl' (with their stomachs on the floor), which itself includes 'inchworm' movements (raising the body onto knees and/or toes in order to push the body forward), leg-only crawls, army crawls, and several other variations. Some infants even avoid crawling in favor of a logroll motion (Karen E. Adolph et al., 1998; Trettien, 1900). Even

so, most infants do start with the iconic crawl. But most infants also employ bear-crawling (crawling on hands and feet), and many (up to 21%) employ bum-shuffling (crawling on hands and buttocks) and crab-crawling (where one foot and/or knee pulls the body forward along the buttock of the opposite, folded leg) (Karen E. Adolph et al., 1998; Fox et al., 2002).

Over several weeks of crawling, infants will develop and maintain a wide repertoire of movements. They tend to continue to add new variants to their repertoires, and to switch between them rather than continuously employing one singular method (Karen E. Adolph et al., 1998; Freedland & Bertenthal, 1994). However, various methods have their own downstream advantages. For instance, the experience of belly-crawling helps make infants' limb movements larger and stronger and gives them a developmental advantage relative to non-belly-crawlers. For instance, when ex-belly-crawlers begin crawling on hands and knees/feet, they are twice as proficient in terms of the size and speed of their movements than their non-belly-crawler counterparts.

However, there are several stages of upright mobility between crawling and walking. Many children begin moving upright with the aid of their parents stooping over and holding their hands above them; some children begin 'walking' on their knees with their hands free in the air; but most infants 'cruise' by moving sideways while supporting themselves with both hands against a support (usually furniture) (Adolph et al., 2011; Vereijken & Adolph, 1999). Here, too, there are some advantages that carry over into later performance. For example, even before they begin to walk, 'cruisers' demonstrate less side-to-side wobble, transfer more of their weight from their hands to their feet, and

are more sensitive to haptic feedback from their hands relative to their non-cruiser counterparts (Haehl et al., 2000; Metcalfe & Clark, 2000; Beatrix Vereijken & Albers, 1998). Indeed, after a period of sideways cruising, many cruisers eventually turn to face forward, supporting themselves with only one hand (Vereijken & Waardenburg, 1996). This brings them that much closer to walking completely independently while also enjoying some of the benefits of their previous two-arm approach.

3.1.4 – Non-linear progress:

The picture in the previous section seems to suggest that progress is relatively linear – one stage cleanly leads into the next and so on. However, despite the various advantages that can come from previous methods, infants must often learn new aspects of performance as they proceed into the next method: each phase brings its own challenges that the infant must learn to navigate. For example, new crawlers are not very good at distinguishing safe from risky ground – many toss themselves over extremely steep slopes and high cliffs during experiments, requiring experimenters to rush in and catch them (Adolph, 2000; Adolph et al., 1997; Bertenthal et al., 1984). However, with time, crawlers learn to understand the limits of their environments and their abilities and so avoid these pitfalls or find safe ways over them (Adolph et al., 1997). Indeed, experienced crawlers can make fine distinctions between navigable and non-navigable surfaces (Adolph et al., 2008). A very similar story is true for cruisers. New cruisers won't have the same success as experienced cruisers, who can distinguish between navigable paths and paths with gaps in their handrail support that they cannot manage (Adolph et al., 2011). This suggests a brief 'plateau', if you will, as the infant enters each

new motor benchmark. Though they do eventually progress and push past it to the next one, each new benchmark requires new lessons to be learned for which the infant starts ‘fresh’ each time.

But this non-linearity extends even further. Not only is it the case that each benchmark brings its own challenges, it also seems that whatever infants learn in precursory benchmarks might not directly transfer to walking. Indeed, walking often requires re-learning how to do various things and applying them differently in this new context. For instance, infants who learned to navigate slopes or cliffs during crawling or cruising often thrust themselves over said slopes or cliffs when they start walking – the previous lessons do not automatically transfer to the new mode of mobility (Adolph et al., 2008, 2011). Furthermore, the learning process in this new case of walking is not quickened by previous experience (Adolph et al., 2008; Adolph et al., 1997).

Even further, learning to walk has more recently been thought to cause a temporary backslide in motor achievements developed in earlier stages. For instance, new walkers often revert back to reaching for small objects with both hands, even while sitting, despite having previously refined their manual control enough to reach with one hand (Corbetta & Bojczyk, 2002). Similarly, some new walkers revert back to less mature postural responses when tested while sitting: this is thought to be because the change to their balance control while upright is somehow interfering with their balance control while sitting (Chen et al., 2007).

All of this suggests a few things. First, that walking, acquired through gradual progression through a series of benchmarks, each requiring its own parameters to be

learned and often a relearning of previous lessons as they apply to the new circumstances, can be understood as a learning process writ large, i.e. Criterion 1. Infants are constantly pushing themselves into new ZPDs, learning about each one and finding its limits, and pushing themselves onto the next one, each of which requires a certain level of mastery in order to do. In this way, we can also see that moving through all these ZPDs requires a level of intentionality and trying to ‘get it right’ (i.e. Criterion 3). To keep pushing themselves in this way, there must be some underlying motivation (which we will discuss more in a later section) that drives them to keep refining their abilities and find the most effective ways of getting about.

What we have also seen, however, is that the acquisition of walking is a process that is constantly setting and resetting various baselines such that learning to walk becomes a much more complicated process than it might seem initially. Rather than there being a singular line of progress, learning to walk relies on several lateral lines of progress, all mingling and intermingling with each other and contributing to the overarching achievement. As the infant moves from one benchmark to another, some contributing lines of development may strengthen while others may weaken – getting to walking requires a full coordination of all these aspects. As such, we can already see how the picture of skill acquisition discussed in earlier chapters already seems a bit overly simplistic.

3.1.5 – Environmental Factors in Walking

But the picture we’ve just established can be made even more complicated. The infant is not learning to walk in a vacuum alone – they find themselves within larger social

systems which influence how they walk. Though the benchmarks of crawling, cruising and walking are common in contemporary Western cultures, they are not universal across cross-cultural and historical data. So, we must not only be concerned with the idiosyncratic ways in which an infant learns to walk, but also with their larger context: a context that includes general social ‘scaffolding’ surround and supporting the infant’s journey to walking.

Beginning with historical variation, even in our own western cultures, the current typical trajectories of benchmarks and onset ages vary from those even 60 years ago. For instance, 40% of late-nineteenth century middle-class American infants skipped crawling and, instead, bum shuffled, crab crawled, spider crawled (on their backs), and logrolled (perhaps because of the fact that infants at the time were clothed in long dresses and so these methods avoided getting their limbs caught in the dress) (Trettien, 1900). Similarly, the average middle-class, American infant in the 1920’s and 30’s began walking between 13-15 months (Gessell et al., 1934; Gessell & Amatruda, 1941; Shirley, 1933). By the 1960s, the average age for walking began just over 12 months (Bayley, 1969) and, today, the average is between 11-12 months (Berger et al., 2007). This continuous shift is thought to be explained by general improvements in nutrition and living conditions (Kuklina et al., 2004).

Cross-cultural variation demonstrates a similar trend. For example, about 29% of Jamaican infants whose mothers used traditional infant massages and exercise regimes skipped crawling all together, while some crawled and walked concurrently (Hopkins &

Westra, 1989, 1990). Similarly, 17% of British infants skipped crawling all together; instead, 10% bumshuffled and 7% just stood up and started walking (Robson, 1984).

But these variations in onset can be pushed even further. Infants from various parts of Africa, and even infants of African descent in North and Central America, have been observed to demonstrate earlier ages of onset for these benchmarks than their non-African counterparts. For example, infants in the Kampala region of Uganda were observed to start standing on their own at 7 months, cruising at 8 months, walking rather ‘stiffly’ at 9 months, walking ‘well’ at 10 months, and running at 12 months (Geber & Dean, 1957). These results have been widely replicated across other parts of Africa and in more recent timeframes (Leiderman et al., 1973; Iloeje et al., 1991; Kilbride et al., 1970; Werner, 1972). Similar accelerated trends have been demonstrated in infants of African descent in the Caribbean (Hopkins & Westra, 1990) and in the United States (Bayley, 1965; Capute et al., 1985).

Much of this acceleration is attributed to differences in childrearing practices. African mothers (and/or grandmothers) engage in daily massage and exercise routines specifically aimed at facilitating their children’s development of walking. They do this using several techniques. Some are done while balancing the infant on a knee or outstretched arm and rubbing infants from head to toe, stretching their arms behind their backs, and/or pulling their knees to their chests. Others involve throwing the infant in the air, holding them by the head, wrist or ankle and shaking them, sitting and standing infants in their lap or on the ground, bouncing them, and generally urging them into upright positions (Adolph, Karasik, et al., 2010; Bril & Sabatier, 1986; Rabain-Jamin &

Wornham, 1993; Super, 1976). Indeed, infants are rarely lied down and are instead carried in slings, even for naps. Even this is thought to encourage walking insofar as it results in additional shaking, rocking and gravity resistance (Bril & Sabatier, 1986; Super, 1976). These techniques have been shown to have a ‘dose-response relation’ with walking – that is, the more frequently and consistently the techniques are applied, the earlier the infant begins to walk (Hopkins & Westra, 1990; Rabain-Jamin & Wornham, 1993; Super, 1976).

But cross-cultural variation cuts the other way, as well – infants whose movements are restricted have delayed onset ages for walking. For example, parents in the rural Shandong and Hebei regions of China commonly use a technique called ‘sandbag rearing’ until their children are potty-trained. On this method, infants from the ages of 12-24 months are kept on their backs in a small bag filled with sand for the majority of the day (between 16-20 hours). The sand acts as a diaper and is changed once a day. For infants reared in this way, 13% can walk at 13-months-old and 72% can walk at 15-months-old. Compare this to 71% and 89%, respectively, for infants of the same age reared in different manners (Xie & Young, 1999).

The kind of social ‘scaffolding’ described in these instances demonstrates how personal development and action, alone, cannot account for the process and manner in which infants learn to walk – they will always be influenced by external ‘scaffolds’ as well. Learning to walk, then, is an even more idiosyncratic process than we initially described: it is one specific to an individual infant in given place and time within a particular social structure.

3.3.1 – Adult ‘Endpoints’ in Walking:

Even though most humans are eventually able to walk, even adults with average bodies across the world demonstrate tremendous variation in their walking. For example, for Western walkers, the energy costs of carrying a heavy load (usually measured by the rate of oxygen consumption relative to ‘normal’ walking conditions) is directly proportional to the weight of the load (Taylor et al., 1980). However, for East African women in the Luo and Kikuyu tribes, they can carry a load weighing 20% of their body weight for ‘free’, i.e. with no increased energetic costs (Heglund et al., 1995; Maloiy et al., 1986; Taylor, 1995). Loads larger than 20% of their body weight do incur a heavier energy expenditure, but not at the same rate as for Western walkers: an increase in load from 20% to 70% creates a linear increase from 0% to 50% energy expenditure, meaning their energy expenditure is continuously offset but the ‘free’ 20% (Maloiy et al., 1986). Similarly, both male and female Nepalese porters carry loads heavier than their bodyweights, while barefoot, up steep, high-altitude paths near Mount Everest. Like the East African women, they can carry 20% of their body weight for free. But with heavier loads they are even more economical: they can carry loads 30% heavier than the maximum weights carried by the East African women for the same energetic costs (Bastien et al., 2005).

These energetic efficiencies are thought to be due to their training and manner of walking: in short, they take advantage of the pendular swinging movement of their alternating legs. During walking, one’s center of mass shifts in several ways: it accelerates and decelerates along a horizontal axis and rises and falls along a vertical

axis. During each step, gravitational potential energy at the center of mass rises along the vertical axis and is then transformed into kinetic energy as it moves forward along the horizontal axis. This process repeats with each step (Carrier et al., 1984; Cavagna et al., 1977). In a perfect pendulum, there would be a 100% energy transfer and no muscular work would be required. For Western walkers, the exchange is at about 65%, and so muscles must step in to do some of the work (Taylor, 1995). However, the East African women and the Nepalese porters can recover more of their potential energy between steps, transferring more than 80% of it (Heglund et al., 1995). In this way, their walking is more efficient than that of Western walkers.

Similarly, people with differently shaped bodies than the average will learn to walk in different ways. For example, people with achondroplasia dwarfism, which results in short stature, a long torso and relatively short arms and legs, learn to walk with a gait that is more ‘bouncing’ and ‘waddling’ than that of the average person. Their initial crawling and walking is delayed compared to the average infant, but they do, of course, learn to walk. When they do, the tilt of their hips is twice that of an average walker, which aids in propulsion and increases the length of their steps (Egginton et al., 2006). Their ankles are extremely flexed to help their long feet clear the forward swing of relatively short legs and, partly because of this excessive flexion, their knees are bent and bowed outwards. Some achondroplastic people continue to walk in this way even after surgical interventions that elongate their legs or even just when wearing stilts (Dominici et al., 2009).

Furthermore, walking will require various adaptations as the body changes, for example, as one gains weight. Obese people, need to use more energy to move their bodies, but they use considerably less than might be expected based on a strict biomechanical calculation (Browning & Kram, 2009). They achieve this by reducing the relative costs of swinging their legs. Higher body mass creates greater impact on the ground which in turn puts greater stress on muscles and joints. To reduce this impact along a ventral dimension, obese people walk slower, increase periods of double support (where both feet are on the ground), and decrease periods of swing (Browning & Kram, 2007). To reduce this impact along a lateral dimension, they widen the side-to-side distance between their feet and swing their legs more to the side than to the front (Abdulrahman & Zebas, 1993).

3.4 – Fluency and Flexibility of Walking:

A&R explicitly point towards fluency and flexibility as core features of skill. A&R characterize *fluency* as follows: it is “what makes skills efficient, coordinated, and beautiful to observe. It is the ability to execute movements smoothly, accurately, and rapidly” (A&R, 2013, 14). Furthermore, A&R note that fluency requires real-time control of movements (ibid, 17). They understand *flexibility*, on the other hand, as “what makes skill truly functional. It is the ability to adapt to changes in local conditions by selecting, modifying, discovering, and creating movements appropriate to the current situation” (ibid). Indeed, A&R emphasize that “variable and novel circumstances are the norm, not the exception”, and so rather than merely reproducing the same behaviors, skill must be able to “modify ongoing movements, select from among a variety of movements in the

repertoire, and generate new forms of movements to cope with a changeable body in a variable world” (ibid, 17). Because of this, they suggest that flexibility involves ‘prospective control of movements’, i.e. preparation and guidance of actions into the future (ibid).

With this in hand, they carefully lay out the literature suggesting that these features are gradually developed in infants’ acquisition of walking: the more practice infants get at various stages of mobility, the more fluent and flexible their behavior in these stages becomes. Our purposes for going through this literature will be very similar: doing so will help demonstrate that infants’ walking does gradually become fluent and flexible in ways that reveal their skillfulness (Criterion 2), but also that their performance is above a contextually appropriate threshold (Criterion 4) and that it is apt (Criterion 6).

3.4.1 – Fluency

New walkers are not yet fluent in their walking. Indeed, because they still have difficulty maintaining balance on one leg while the other swings forward, they take steps that are ‘small, jerky, and uneven’ (Bril and Berniere, 1993). This is thought to happen for several reasons. First, in these early stages, they spend 40-80% of their gait cycle with both feet on the ground and only very short periods when one leg is swinging forward (Ledebt & Bril, 2000; B. Vereijken et al., 1998). This means that they are generally slow in their movements, which means they don’t have enough time to flex their ankles backward while they swing their leg, which, in turn, means they tend to land either flat-footed or on their toes. Similarly, they hold their hips and knees in a tight, flexed position, and so neither cushion the impact of footfall nor fully extend their leg at toe-off

(D. Sutherland et al., 1980). Lastly, muscle action in the leg is so inconsistent as to often be in conflict with itself (i.e. flexor and extensor muscle are contracted simultaneously) (Chang et al., 2006). All of this means that their heads and trunks are quite wobbly which confounds their balance even further (Bril & Ledebt, 1998; Ledebt & Bril, 2000).

To compensate for all of this, infants take short steps, with legs splayed far apart, often with a wider distance between their feet than they take between each step forward (Karen E. Adolph et al., 2003; Bril & Brenière, 1992). Indeed, the timing and distance between steps are asymmetrical and variable, which suggests that infants are constantly trying to regain their balance between steps (J. E. Clark et al., 1988).

To deal with these issues, infants adopt a wide variety of means to improve their balance and forward progression. They will usually explore all of their space and try as many configurations of their body as possible, utilizing whatever seems most effective (Adolph et al., 2003; Bril & Brenière, 1992; Bril & Ledebt, 1998; Ledebt & Bril, 2000; McGraw, 1945) All in all, infants will look for ways to “exploit – rather than fight – the gravitational forces acting on their limbs and the inertial forces generated by their limb movements” – in short, they look for ways to “do less to get more” (A&R, 2013, 16).

Eventually, ‘double-support’ periods (i.e. those in which both feet are on the ground) diminish to 20% of the gait cycle and swing periods increase until they reach near adult ratios (Shumway-Cook & Woollacott, 2012). Overall speed increases and they get enough time to flex their ankle and pull their toes up during swing so that they can land on their heels. They relax their hips and knees and start keeping their legs closer together (Adolph et al., 2003). Muscle action becomes more reciprocal, and they are able

to stabilize the wobbling in their head and bodies (Assaiante, 1998; Ledebt & Bril, 2000). As this happens, their performance demonstrates significantly less variability and their fluency increases. In general, infants reach this more fluent state within 3-6 months of having started walking (Adolph et al., 2003; McGraw, 1945), but don't usually reach adult levels of fluency till 5-7 years old (Bril & Ledebt, 1998; D. H. Sutherland et al., 1988).

3.4.2 – Flexibility:

Infants learning to walk must be constantly responding to shifts in their bodies, skills and environments. For example, walking while making even subtle shifts of the head or arms, or even taking a deep breath, will require a shift on the infant's part to maintain balance. Similarly, as the infant grows and their bodily dimensions change, so will their center of gravity. Furthermore, as they move through various benchmarks of mobility, they will encounter new objects and obstacles that they will have to navigate. For example, as an infant learns to crawl, they will come in contact with new objects, or familiar objects in new ways than they did while they were stationary. The same goes for when they start walking (Klopfer, 1988). All of this means that walking requires a great deal of flexibility.

To test an infant's flexibility in movement, researchers will introduce them to novel challenges to movement: barriers to step over or duck under (Schmuckler, 1996; Van der Meer, 1997), openings to squeeze through, waterbeds and foam pits (Gibson et al., 1987; Joh & Adolph, 2006), slippery shoes and high-friction surfaces (Adolph, Joh, et al., 2010) and more. However, the classic paradigmatic test for flexibility is the 'visual

cliff': a glass table with a patterned surface just below the glass on the 'shallow' end and then the visible floor on the 'deep' end (Gibson & Walk, 1960). Infants are encouraged to traverse this table by a parent on the other end beckoning to them. Note that the drop-off is only an illusion, since the infants are on a glass table that extends over and across the 'cliff' – even so, it is invisible to them and they are not aware of this. Infants new to crawling or walking traverse the table with no hesitation, while more experienced infants avoid what looks to them to be a drop-off (B. Bertenthal et al., 1984; B. I. Bertenthal & Campos, 1984). That is, infants who are new to crawling and walking have not yet gained the flexibility of adjusting their movements to the circumstances at hand: they simply move the way they usually do regardless of what they encounter. However, more experienced crawlers and walkers have come to recognize that different circumstances will require different applications of their mobility and are able to execute this, demonstrating a level of flexibility not seen in the previous instance.

However, a methodological problem is raised in these experiments: through feeling and at least one experience over the 'deep' side, these more experienced infants eventually come to understand that the glass extends across the table and so learn the 'trick'. After this, they begin to cross normally and no longer attempt to prepare for the 'cliff' (Campos et al., 1978). Because of this, infants can't be tested more than once. Yet, some methods have been devised to get around this issue: these include introducing actual cliffs, gaps, or slopes in the surface of a table, etc. Since the cliffs and gaps are real, in these instances, researchers must follow alongside the infants as they traverse these paths to keep them safe: even so, the researchers don't catch the infants or stop

them from falling in most cases, so they can observe the same infant attempting to cross several times and utilize different approaches.

Similar results are demonstrated in these ‘real’ cases, too: amateur crawlers and walkers coolly attempt to cross risky cliffs and gaps. For instance, inexperienced crawlers may gaze down a slope with both hands touching the angled surface and then proceed to plunge down the slope headfirst (Karen E Adolph et al., 1997). Similarly, novice cruisers may dangle their foot over a gap for several seconds then tip themselves into the empty space and fall (Karen E. Adolph et al., 2011). New walkers stop at the edge of the cliff, look down, and then step right over it.

However, after a few months of crawling or walking, infants demonstrate careful selection of actions to navigate the same obstacles. For instance, while on an easy path, they may crawl or walk fluidly, without breaking stride. But, on a more difficult path, they will slow down and modify their gait to carefully get around the obstacle. On especially risky paths, they might refuse to cross altogether, but often find different ways to get around the obstacles such as turning around and backing down the cliff feet first, etc. In general, when infants recognize that a path is risky, they will begin by gathering more information about it, either from tactile exploration or constantly checking in with visual feedback: they may stretch their arms over the gap to see how wide it is, they may stand with their feet at the edge and rock back and forth (apparently to generate torque around their ankles), and all this tactile exploration is usually accompanied by concerted looking and vigilance (A&R, 2013).

Other methods besides the visual cliff experiments have proven especially useful in testing for infant flexibility in mobility: namely, experiments using adjustable slopes and bridges. These are particularly useful for several reasons: the angle of a slope and width of the bridge can be adjusted to allow more precise assessment of accuracy in performance; slopes and bridges can be navigated in a large number of ways and so allow the researchers to observe the infants' existing and developing repertoires of solutions; and, because getting across these paths will require multiple steps, these set-ups gives researchers the opportunity to study gait modifications.

Again, bridge and slope experiments have yielded results that largely align with what we've already seen but demonstrate the degree of sensitivity in infant's choices of action. Experienced infants can traverse steep slopes and narrow bridges by adjusting their walking gait, usually by shortening the length of their steps and slowing down (Adolph & Avolio, 2000; Berger et al., 2005, 2010; Berger & Adolph, 2003; Gill et al., 2009). In one instance, an infant made their way down a 35-degree slope in 22 steps, compared to walking the same distance at a shallow 6-degree slope in 5 steps. Even changes as small as 2 degrees on the angle of a slope and/or 2cm change in the width of the bridge can have impacts on the infant's gait.

Furthermore, results demonstrate the prospective nature of the actions – plots show tight clusters of steps at the edge of an obstacle, demonstrating a shortening of gait as they approach, and then maintenance of small steps as they traverse it. Similarly, infants brace between steps while going down a steep slope to prevent themselves from falling forward, and they turn their bodies sideways, leading with one leg, as they go

across narrow bridges. In all these instances, it seems the infants recognize that a given stance will be more beneficial before they even get started, as opposed to making adjustments mid-way through the course.

However, A&R identify the following as the strongest proof of flexibility in infants' mobility: namely, "their spontaneous search for new means and the retention of a variety of means in their repertoires" (A&R, 2013, 19). In general, infants seem to be constantly looking for ways to get around whatever obstacles they encounter and develop new means of doing so either by developing new means altogether or recombining old ones in new ways. For example, on extremely difficult steep slopes, infants will try several different means to get down. In many cases, infants may start by sitting at the edge of the platform and trying to turn around and back down feet first, but then get back up to standing position as if to 'reassess' from safe ground, and try something else (Karen E. Adolph, 1995; Karen E Adolph et al., 1997). Furthermore, when they discover a new method, they may delightedly repeat it through several trials, but soon enough they will start trying to find another new way. Over the course of a session, infants are likely to try several ways of navigating the obstacles. However, they will often exhaust all the possibilities of one method before moving on to another. Often, they maintain given strategies over a period of a few weeks or months, but their use will gradually decrease as new or more efficient methods are developed. Even so, the older methods never disappear completely.

3.4.3 – Upshots of Empirical F&F studies:

All in all, we have seen compelling evidence to suggest that infants begin various stages of mobility, especially crawling and walking, without a great deal of either fluency or flexibility. As novitiates in these stages, their actions are not executed “smoothly, accurately, and rapidly”, i.e. fluently, nor do they demonstrate a dexterity in “selecting, modifying, discovering, and creating movements appropriate to the current situation”, i.e. flexibility. Indeed, their walking is jerky and halting and is not responsive to cliffs, slopes, or other obstacles. However, as their skill grows, their movement does eventually become ‘smooth, accurate and rapid’, especially in normal circumstances, and becomes varied and responsive enough to be able to adapt to various obstacles.

Indeed, their movement can be assessed under the RRRR criteria. Starting with *resilience*, we can see quite clearly that their walking in early stages does not meet this requirement – it is rather weak, even in normal circumstances, and cannot be maintained across other circumstances. A very similar picture can be painted for *reliability*. When infants start in various benchmarks, their abilities are not totally reliable and, even in normal circumstances, may be rather unwieldly. Indeed, evidence suggested that their balance was so poor in these early stages as to not be reliable even then. Similarly, their walking is not yet *responsive*, i.e. adjustments relative to the environment and specific goings on of an event are not yet being accounted for. Because of all of this, new walkers’ skill is not yet *robust*.

However, with time, experienced infants’ walking does meet the RRRR requirements. For instance, it does meet the *resilience* criterion, i.e. infants are able to be

walk (or be otherwise mobile) in the face of various obstacles and can maintain a relatively high level of performance. Similarly, their walking becomes more *reliable* – they can walk relatively well across many different circumstances in the face of various obstacles. Lastly, their walking demonstrates a *responsiveness* to circumstance, even a preemptory one: infants will make adjustments to their walking in response to circumstance and can even be shown to select appropriate modes of moving in order to preempt foreseen difficulties of a circumstance. Because of all of this, their walking skills can be understood to be *robust*.

Along the same lines, I suggest that, because experienced infants' walking is robust according to the criteria discussed above, it should also be understood to be performance above a context specific threshold. Recall that infants don't usually reach the same level of fluency as adults until 5-7 years of age, so the threshold for what constitutes skilled walking for infants and young children should not be understood as the same as that for adults. Long before they reach adult levels, one would usually feel confident in saying that a young child walks as well as anyone else. Even further, recall that various environmental factors at play in the age of onset for various mobility benchmarks suggest that there might not even be a true threshold relative to age. That is, judging the threshold relative to how well an infant walks for an infant of their age might seem culturally specific and only sets a relative threshold to hold them against. However, experienced infants who meet the RRRR criteria are demonstrating a much higher level of skill than they had previously and suggests that their performance could be reasonably understood as sufficient for skillfulness.

Lastly, given the robustness of their walking, it seems clear that their success at movement at later stages of acquisition is apt, i.e. due to their own acquired skill rather than to any accident or luck. Initially, infants fail because they have not yet acquired the relevant skill(s). But as they gain experience, they begin to succeed precisely because they have acquired the right skill(s). Because their walking is learned, rather than innate, they are truly succeeding because of their own skill gained over days, weeks and months of practice. Similarly, they are not receiving guidance, tips or hints from researchers, nor are they intervening in the infant's performance (except to catch them in particularly risky circumstances). Infants' success at walking in general, as well as navigating all these various obstacles, is due to their own practice and acquired experience. As such, it is apt.

3.5 – Knowledge and Walking

The conversation knowledge and skill in the previous chapter focused on questions about what kind of knowledge was at play, relations between different kinds of knowledge and skill, as well as priority of either knowledge or skill. However, developmental psychologists have a different focus when discussed knowledge and walking, in particular. They often focused on walking as a catalyst for development in other domains of growth (B. Bertenthal et al., 1984; Campos & Anderson, 2000; Gibson & Pick, 2000; Mahler et al., 1975; Piaget, 1952; Thelen & Smith, 1996). Indeed, as Eleanor Gibson puts, walking sets off “a kind of cognitive revolution”, one that “must result when an infant's horizons are expanded by their acquisition of self-initiated, self-controlled locomotion. A new field of knowledge is opened up and a whole new set of

skills must be mastered. A new kind of activity that is both exploratory and performatory becomes available for learning about the larger world” (Gibson, 1988, p. 27). Along the same lines, others include additional motor developments, perception, affect, social behavior and social understanding as downstream effects of learning to walk. In this way, developmental psychologists tend to see the acquisition of the skill of walking as a pivotal moment for the acquisition of knowledge more generally: in the process of acquiring the skill of walking, more and more domains of knowledge are opened up and are intimately tied to direct experience in this process.

Notice that a key object of interest in Gibson’s thinking is that self-initiated, self-controlled movement – this seems to be the lynchpin in the changes she describes. This kind of direct experience has been shown to be critical for the kinds of advancements he portrays. In principle, being passively carted around and carried by parents could expose children to many of the same things as they would get from their own movement: they would see the same things, see others interacting with them, etc. But practically, this is not the case: much like the passive kitten in (Held & Hein, 1963) classic ‘kitty carousel’ experiment, pre-locomotor infants do not perform ‘adaptively’ to various motor tasks (Campos et al., 1992). This only begins to happen when they start moving around, themselves.

Being able to move around oneself opens up the possibility of learning about one’s own body, others’ bodies, objects, and larger features of the environment in ways one simply cannot do when immobile. Indeed, they begin to learn about the kinds of movement their bodies are capable of, what they can and can’t support, how to move

between ZPD's, start to recognize things they can imitate from other people, etc.

Moreover, as they move around an environment, infants will not only learn about their own and other bodies, but also encounter new opportunities for crawling, climbing, walking, etc. which will, in turn, open up possibilities learning about which surfaces can bear their weight, for new obstacles, falling, etc. (Gibson, 1988). As an infant learns to crawl, for example, they will encounter new areas and surfaces and begin to learn how to move over those. Similarly with climbing, cruising and walking, as they move about, they will encounter new areas and begin to find ways to navigate them.

Even further, however, is the fact that moving around an environment opens up spatio-temporal dimensions of learning – tying together environments based on past and current locations of the body and relative to landmarks or previous places visited (Bushnell et al., 1995). Self-driven movement (such as crawling, walking, etc.) has been shown to improve infant's performance on spatial search tasks (B. Bertenthal et al., 1984; Kermoian & Campos, 1988) and on place learning and cue learning tasks (M. W. Clearfield, 2004). It also helps infants remember locations as compared to being carried around by their parents (Acredolo et al., 1984; Benson & Užgiris, 1985). In general, self-generated mobility encourages infants to attend to a larger array of location cues more efficiently. It also encourages them to weigh different types of information (for example, bodily-based cues or landmark-based cues) and increases their efficiency and adaptability in these regards (Newcombe & Huttenlocher, 2000).

Further still, becoming independently mobile opens up new motivations for movement and further exploration. As infants start to move around, they become more

interested in objects at a distance - they begin to recognize that it is within their potential reach insofar as they can bring themselves to that object and even carry it or transport it. All of this encourages even more visual, motor and manual exploration of objects -the beginning of a chain reaction, as it were (Melissa W. Clearfield et al., 2008; Gustafson, 1984; Karasik et al., 2011). Indeed, carrying becomes an important new domain of focus, with bouts of carrying increasing by 520% over the transition from crawling to walking (Karasik et al., 2011). Similarly, carrying helps facilitate social motivations and make them important, with infants taking objects to adults to share with them, rather than waiting for adults to come to them (Karasik et al., 2011).

Indeed, walking seems to open up a new phase of social agency for infants: though crawlers can move away from their caregivers and explore their surroundings, walking is thought to be an important point of psychological separation from caregivers and starts infants' 'love affair with the world' (Mahler et al. 1975). During this phase, for instance, parents report a greater testing of wills than had previously been displayed: both infants and parents increase their frequency of saying 'no' (Biringen et al., 1995, 2008), and parents are observed to utilize different language, gestures and displays of facial emotion between crawling and walking infants (Campos & Anderson, 2000; Karasik et al., 2008; Tamis-LeMonda et al., 2007).

All of this suggests a different approach to the relationship between skill and knowledge than we had previously explored. Rather than focusing on what kind of knowledge is in play at what times and which is prior to which, etc., the picture here can be understood as a mutually reinforcing one, with skill opening the path to knowledge

which opens the path to new skills which opens the path to new knowledge, and so on and so forth. Of course, the question of whether know-how or know-that is at play, for instance, can still be asked. But perhaps, for the purposes of a conversation around skill acquisition this kind of question is not of primary importance. Indeed, even for a conversation around the nature of skill itself, it may be more productive to focus on the mutually developing and reinforcing ‘feedback-loop’ between skill and knowledge than what specific kind of knowledge is at play. Asking our original kinds of questions doesn’t help us understand the ways in which direct engagement in an environment contribute to acquiring new knowledge or understanding, or how new goals and motivations can be introduced as development progresses, etc.

3.6 – Developmental Acquisition of Skill:

It seems fair to say that the developmental metaphor for skill that was in use in the previous chapters was like that of an evolutionary chain, or a roughly linear series of milestones until achievement of skill proper. Though there may be some disagreement amongst various camps about what the specific links in the chain/milestones are, what they consist in, and whether there is a true end to the process, the general picture seems to be similar: one starts as a novice and, gradually with practice and experience, achieves increasing levels of mastery until one is a true master. The developmental literature we’ve just explored, however, gives us some reason to push against this simple picture.

A&R consider another metaphor within developmental psychology that has taken root over the last half century or so; namely, the developmental cascade. On this picture, most famously put forward by (Waddington, 1957), development is like the movement of

a ball down an epigenetic landscape. Picture a landscape with many hills, bumps, valleys, furrows, etc.: the organism's development, pictured by a ball which starts at the top of the hilly landscape, will gradually move down in accordance with the contours of the environment it encounters, but only probabilistically. Indeed, small influences at the beginning of the journey can have large impacts as the journey proceeds. Even so, there seems to be a general convergence across the landscape to a few larger 'canals', with many of the smaller trails and valleys leading into them as tributaries. As such, development can be understood to follow a few key routes down to the bottom of the landscape, usually because of self-correction along the way.

Similarly, (Thelen, 2005) uses the metaphor of a mountain stream; a continuous, ever-flowing path of activity and change, all influenced by the underlying streambed. This stream is continuously influenced by the landscape it runs through, the conditions of the streambed, the weather surrounding it, etc., and so carries its history with it at all times. Different real-time actions will have different effects on the stream depending on its immediate conditions (i.e. throwing a rock into the stream will have different impacts depending on where along the stream one is, the conditions of the water, etc.). And the progress of the stream has both short- and long-term down-stream effects: it can support various lifeforms, it may create runoff, it may carry pollutants, it may contribute to groundwater sources, etc. In this way, the stream creates history, as well.

However, A&R push against these kinds of models, noting that they don't account for the child's role, or agency, in the process of development. In this sense, they go too far in the direction of environmental, 'external' contributing factors. Yet, there are a

number of views known as ‘infant as’ views – infant as blank slate, infant as mini-computer, infant as scientist in the crib – which do highlight the infant’s role in development. Even so, A&R argue that this latter set tends to lose track of the larger developmental story and so go too far in the ‘intentionality’ direction (e.g. (Gopnik et al., 1999).

Aiming to try to capture a middle-ground view, A&R offer another model: that of a root or branch from a tree, or, more specifically, a banyan tree. In general, when a root or branch starts growing, it is relatively weak but flexible and gradually begins to extend into its environment. As it does this, it encounters various obstacles and may change its direction of growth around or through them. This means that each root or branch “can develop its highly individualized and idiosyncratic form, while adhering to the defining features and functions of its kind” (A&R, 2013, 25).

Banyan trees, in particular, serve as a useful example, because the growth of their branches, roots, and trunks are deeply intermingled and merge in and out of each other. Banyans start as epiphytes⁴¹ growing on parent trees. They initially derive moisture from the air and rain until their roots, growing from their branches, grow long enough to reach the ground. These aerial roots eventually form new trunks, with the roots and branches diverging and merging into each other. Because each tree faces unique circumstances and environments, and will respond accordingly, banyans can grow in an extremely wide variety of ways and assume many different sizes and shapes. This makes simple categorization of developmental milestones for banyans a very difficult task.

⁴¹ A plant that grows on another plant but isn’t a parasite.

Understanding skill in general in this way has a number of advantages. First, we can maintain a general picture of ‘normal’ or ‘common’ steps or levels one might traverse as a skill is acquired and developed, but still leave plenty of room for individual idiosyncrasies and unique experiences along the way. Second, it opens a way for including both ‘environmental’ factors as well as the subject’s own intentionality into the process, and showing their interconnected, dynamic development in response to each other. Lastly, it gives us a new of understanding how one development can lay the groundwork for another in a much more nuanced way. All of this was lacking from our previous discussions and will prove useful as we move forward.

Conclusion:

I have demonstrated that walking meets all the criteria of skill we had laid out in previous chapters and, as such, should be understood as a skill. I have also recommended subscribing to a new metaphor of skill which includes much of the lessons gleaned from the developmental psychological literature around walking. This serves to fill in some gaps from our previous analyses and helps set us up for further exploration in a much more careful and nuanced way.

I set out the example of walking as a skill because I believe it embodies the vast majority of our skills: ‘basic’ skills that help us get through the world ‘normally’ on a day to day basis; skills that we may not normally think of as skills but follow the same general patterns and which demonstrate an interdependence between ourselves and our environments. As we move into the main thrust of the dissertation, i.e. memory of skill, *this* is the kind of skill I want to be focused on, primarily. I will, at times, discuss more

traditional examples of skill, in particular playing sport. However, I will suggest that these more basic skills are at the root of the more traditional ones and that they are the more interesting ones for us to focus on within a conversation around memory. For this reason, it was important to establish something like walking as a skill.

So now, having established some criteria for skill and discussed an example of a basic skill at length, we are ready to move on to the main show: memory and skill. Remember that we began with a brief discussion of the nature of procedural memory – in the case of something like walking, we have just seen how it is that infants gradually learn and acquire the ability to walk. Whatever information and experience this is must be, in some sense, remembered or retained. But whatever this is does not fall neatly under the auspices of facts, propositions, etc. or even past events in one's life. But the question is, how should we model this kind of memory? How are we to understand the kind of memory at play when infants learn and remember how to walk? As we shall see, there are a few standard models of memory (discussed in chapters 4 and 5), which we might be tempted to adopt. But, as we work through these models, we will see that they do not fit or uphold our criteria and understanding of skill. Ultimately, we will need to offer a new model. But, before we get ahead of ourselves, let us then begin to examine these standard models and see more carefully how they would aim to account for the memory involved in skills.

Part 2: **Memory and Skill**

While the first part of this project aimed at getting a clearer understanding of skill (especially as the content of memory of skill), this second part is aimed more explicitly at memory. In particular, the next few chapters aim at examining dominant understandings of memory and determining whether and to what extent they apply in the case of procedural memory, or memory of skill. Chapter 4 targets one dominant understanding, namely Storage and Retrieval. This is a very wide and varied category of theories that has had many iterations over time. However, I group these theories together, and identify a key modern example of it, in order to examine the core underlying assumptions and attempt to apply them in the case of procedural memory. By identifying and examining the insufficiencies of a strict Storage and Retrieval model for understanding procedural memory, we then find the way towards another potential candidate, namely a (re)construction model, and its descendant, the simulation model, which will be carefully examined in Chapter 5. I argue that, though these models are an important modification of classic Storage and Retrieval models, the reconstructive account ultimately maintains some of its core components and falls prey to similar failures. However, with this demonstration, we will be able to see some benefits of their approaches for understanding procedural memory and how we can better describe that phenomenon than some leading understandings. Ultimately, this analysis paves the way for a final account of procedural memory in the last part of the dissertation.

Interlude - Metaphor and Memory:

As we enter the conversation around memory, I want to highlight the role of metaphor in our understanding of memory systems. As we will see, metaphors like Storage and Retrieval and Reconstruction, shape our understanding of what is happening when one remembers. I do this because it seems to me that the use of these metaphors in explanation is thought to be innocuous and that they could be understood merely as loose guides, pedagogical or heuristic, to understanding the nature of memory. However, as these metaphors become more entrenched and repeated, I believe they are continuously reinforced in such a way as to no longer be understood as pedagogical or heuristic, but the true nature of the phenomenon. In analyzing these models of memory, I also aim to assess the manner in which their grounding metaphors limit or distort our understanding of procedural memory and memory at large. Let us, then, begin with a brief exploration of the role of metaphor in language in general, as well as the way it has played out in memory discourses.

Lakoff and Johnson – Metaphorical Structuring of Concepts:

In their seminal book, *Metaphors We Live By*, (Lakoff & Johnson, 1980) investigate the role of metaphor in language and of our general understandings of the world. They begin with the central thesis that, rather than being a poetic or artistic flourish, metaphor is pervasive throughout everyday speech and writing. Indeed, they suggest that a metaphor like “ARGUMENT IS WAR” creates metaphorical entailments which “can characterize a coherent system of metaphorical concepts and a corresponding coherent system of metaphorical expressions for those concepts” (ibid, 9). That is, in this

case, the ‘warlike’ nature of argumentation is deeply and systematically inscribed in our understanding of it and works its way into many subcategorizations of the concept of an ‘argument’. For instance, each party in an argument sees themselves as having “something to win and something to lose, territory to establish and territory to defend. In a no-holds-barred argument, you attack, defend, counterattack, etc., using whatever verbal means you have at your disposal – intimidation, threat, invoking authority, insult, belittling, challenging authority, evading issues, bargaining, flattering and even trying to give ‘rational reasons’” (L&J 62)⁴².

An important feature of any metaphor is, as Lakoff and Johnson suggest, that their use both highlights some aspects of a concept and hides or obscures others. For instance, if ARGUMENT IS WAR, it is much easier to see the ways in which one must *attack* one’s *opponent* or their point of view (or, similarly, put up a *defense* against one’s *opponent’s attacks*). At the same time, this metaphorical framing makes it harder to see elements of cooperation in argumentation and how the two parties might be sharing their time and effort with each other to learn. In this way, metaphors can partially structure a concept and be extended in some ways and not others.

Indeed, L&J explain that the ‘unused’ parts of the metaphor can often be the basis for “figurative” or “imaginative” language about the topic in question. For instance, if

⁴² Here, L&J are talking about a general sense of argumentation and proceed to highlight a more refined version of arguments, RATIONAL ARGUMENTS, which are more common in intellectual and academic circles. However, they suggest that RATIONAL ARGUMENTS generally maintain the same ‘warlike’ structure despite being presented in a more ‘civilized’ and ‘unbiased’ manner. That is, though a RATIONAL ARGUMENT is meant to be carried out on the basis of “stating of premises, the citing of supporting evidence, and the drawing of logical conclusions”, it often maintains, “in hidden form”, all the usual tactics of intimidation, threat, drawing on authority, insult, bargaining, flattery, etc. that are used in order to prove one’s point and win the argument (L&J 63-64).

THEORIES ARE BUILDINGS and the ‘used’ parts of the metaphor deal with *constructions* and *foundations*, some ‘unused’ parts might be the *rooms* or *hallways* inside the building or the *gargoyles* on its façade. So, if someone’s theory is described as ‘having thousands of little *rooms* and long, winding *corridors*’, or as “a massive *Gothic* theory covered in *gargoyles*”, the descriptions can still be understood to operate within the larger metaphor but fall outside the domain of normal literal language. Most metaphors are not literal, and this is because only parts of them are being used to structure a given concept (L&J 53-54).

In general, the strength or success of a metaphor depends on its fidelity to some experiential basis. Lakoff and Johnson suggest that, often enough, language puts together large scale orientational metaphors which organize a whole system of metaphors, usually around an element of spatiality that aligns with physical and cultural experience. Indeed, L&J argue that spatial orientations (such as UP-DOWN, FRONT-BACK, IN-OUT, NEAR-FAR, etc.) are central to our embodied experiences in the world and, as such, are often given priority over other potential spatial organizations (like Cartesian coordinates, for instance). This physical grounding is what makes concepts such as UP important in our conceptualization of other experiences. Social, emotional or cultural correlations with our bodily spatial orientations (such as HAPPINESS being associated with standing up-right, or SADNESS with slouching) can then become the basis for orientational metaphors like HAPPY/GOOD IS UP or SAD/BAD IS DOWN.⁴³

⁴³ However, L&J are clear in explaining that this should not be understood as an argument experiential priority, but rather a statement about the way in which we conceptualize experiences. They say that “we typically conceptualize the nonphysical *in terms of* the physical – that is, we

These orientational metaphors will often draw together varieties of experiences around the same organizing idea – that is, rather than there being many different kinds of ‘up’, there are many different experiences that draw on the notion of ‘verticality’ (as seen in similar metaphors like CONSCIOUSNESS IS UP, HEALTH IS UP, MORE IS UP, STATUS IS UP, etc.). As such, these orientational metaphors can often dominate the manner in which we think about a whole system of extended concepts. Usually, which concepts are drawn together and with which orientations will be chosen from among many possibilities based on the widest range of coherence in a given cultural framework⁴⁴.

However, orientational metaphors can only get us so far. Experiences of physical objects open up another realm of understanding which can then be used to metaphorically understand other experiences. That is, we can create *ontological* metaphors, based on experiences of objects, that give us a way of referring to, categorizing, grouping and quantifying many other experiences. Indeed, metaphorical structures based on experiences of objects give us a means of viewing not only unbounded physical phenomenon, but also events, activities, emotions, and ideas as discrete entities and substances.

For instance, the ontological metaphor THE MIND IS AN ENTITY is based on the notion that we can bound and describe such an abstract thing in the ways we do

conceptualize the less clearly delineated in terms of the more clearly delineated” (L&J 59). In this way, we can equally basic kinds of experiences without having conceptualizations of them that are equally basic.

⁴⁴ This also means that these orientations can vary from culture to culture but will be similarly dominant within each one.

physical objects. Furthermore, this kind of ontological metaphor can be elaborated in various ways, giving us a better grip on the entity. Consider the following two elaborative metaphors of the main metaphor: THE MIND IS A MACHINE and THE MIND IS A BRITTLE OBJECT. The former allows us to talk about how the mind is *operating*, may be *a little rusty*, how its *wheels are turning*, or is *running out of steam*; whereas the latter allows us to talk about its *fragility*, its *need for care*, how it can *break* or be *shattered*, *snap* or *go to pieces*. As L&J explain, “these metaphors specify different kinds of objects” which “give us different metaphorical models of what the mind is and thereby allow us to focus on different aspects of mental experience” (L&J 28). In other words, the ways in which we develop or elaborate on the main ontological metaphors allows us to talk about different aspects or understandings of the same ‘bounded’ entity.

But ontological metaphors can be pushed even further. Once something is understood in physical terms, it can be personified or specified as being a person. This “allows us to comprehend a wide variety of experiences with nonhuman entities in terms of human motivations, characteristics, and activities” (L&J 33). For instance, a THEORY can *explain* or *tell* us something, a FACT can *argue* something, an EXPERIMENT can *give birth to* a new paradigm, a CONSEQUENCE can *catch up with* you. However, these examples not only give us an indication of how an entity can be understood as a person, but also what type of a person. For instance, in our examples, the metaphors are not merely that A THEORY IS A PERSON, or A FACT IS A PERSON;

but rather A THEORY IS A TEACHER or A FACT IS AN ADVERSARY. This not only gives us a way of understanding the entity but also a way of acting towards it⁴⁵.

Taking all of this together, we can see how, beyond merely drawing from experience, metaphors can also give us means of structuring our experiences. Take, for instance, a case of starting a conversation with someone which gradually morphs into an argument. Being able to understand a shift in tone from a CONVERSATION to an ARGUMENT demonstrates several things: first, that a subject is likely using these concepts to help shape their understanding of the events in question and, second, using them to determine how to act and react in response to the change in situation. But how is this shift between concepts being noticed and taken up?

To begin with, L&J suggest that our accumulated experiences contribute to the development of *gestalts*. As L&J understand it, a gestalt is a structured whole, a “complex of properties occurring together” within our experience that is both basic to our experience of the concept and of our experience at large (L&J 71, 117). Noticing the shift from CONVERSATION to ARGUMENTS requires recognizing a shift in experiential *gestalt*. In this case, then, the activity of talking, experienced initially as a gestalt of

⁴⁵ However, L&J are careful to point out that in some cases of personification one is not actually referring to the object in question as a person but is rather using them to refer to something or someone else. For instance, a server at a restaurant may say that “*Table 2* is waiting for their check” – the server doesn’t mean that Table 2 is a person but is using it to refer to the customer sitting at table 2. This is a case of *metonymy* and is distinct from metaphor: metaphor is primarily a means of conceiving of one thing in terms of another while metonymy is primarily a means of reference where one thing can stand in for another. However, metonymies are just as systematic as metaphors, and can help us understand the objects in question in similar ways: whatever the metonymy chooses to highlight about the object also tells us something about how they’re thinking about what its referring to. As such, they can structure our thoughts, attitudes and actions in similar ways to metaphors. Indeed, L&J suggest that metonymy can be a critical link between our everyday experiences and the coherence of metaphorical systems.

CONVERSATION, becomes further structured by a gestalt of WAR. Each of these gestalts is multidimensional and covers many aspects and levels of the concepts in question. But often, they will have some shared elements which helps draw a throughline between them. In this case, ARGUMENT can be understood as a variety of CONVERSATION which shares many core features and behaviors (for instance, two parties/speakers; talking alternately; one person finishing results in the other taking their turn; appropriate beginning, middle, and end; etc.) but refines some to a confrontational context (two adversaries; attack followed by retreat, defense or counterattack; attack, battle, then truce/stalemate/surrender/victory, etc.).

Indeed, L&J suggest that organizing our experiences in terms of multidimensional gestalts makes them into *structured wholes*, which, in turn, makes our experience *coherent* (81). That is, being able to categorize diverse experiences according to some shared dimensions across gestalts is what allows us to see them as intelligibly related: in this case, rather than understanding this one interaction as two separate, or unrelated instances of verbal communication, being able to identify the similarities between the two gestalts, as well as their differences, allows one to understand and make use of the experiential shift in a way that also provides a playbook for how to engage.

For our purposes, then, we should keep in mind how the various metaphors of memory we will discuss in the coming sections and chapters influence our understanding of memory. If L&J are right, then whatever ontological metaphors we use will likely be systematically influencing our thinking around the concept of MEMORY, any of its subcategories, and any adjacent notions. As was initially mentioned in the introduction,

many people might suggest that the use of metaphors is either a basic means of linguistically describing some phenomenon or merely a tool, pedagogical or heuristic. That is, that the use of a metaphor in exploration, investigation, explanation or the like is merely a means of presentation that eases understanding but does not seep down to the level of our ‘true’ understanding of whatever the metaphor describes – there is some pristine concept behind the metaphor that the metaphor does not truly impact. However, I suggest that L&J’s work plausibly suggests that this is not the case: that in fact, the metaphors we use to describe memory deeply influence our understanding of it and its functioning in ways that are difficult to separate from the concept itself. Indeed, the metaphors we use to frame memory are often so deep that they impact how hypotheses are formed, which of those are investigated, and which relationships and findings are taken up in research. In these cases, the metaphors are not mere tools, but become definite of the concepts they are used to describe and create feedback loops in action that reinforce them.

That being said, I do not think that there is anything inherently wrong with the use of, or reliance on, metaphors in research of this kind. Indeed, their use might be inevitable and could, under the right circumstances, be helpful. However, I do think that the ‘dogmatic’ or unconsidered use of metaphors or, similarly, use in such a way that does not allow room for alternative metaphors that may better encapsulate evidence and analysis, is deeply problematic. Researchers should be able and willing to consider alternatives to their general metaphorical narrative and should be willing to investigate the ways in which their adherence to a given metaphor impacts and limits the scope of

their research. Insofar as metaphors are used with this kind of consideration, there is nothing wrong with their use.

However, as we will see in the next section, some metaphors for memory are ubiquitous and, even with a large degree of variation, largely fall within one culturally accepted and intuitive ‘core’ model. As L&J explain, whatever metaphors we use will highlight certain aspects of MEMORY and obscure others. Though the various iterations of the ‘core’ metaphor highlight different aspects of the metaphor, draw on different ‘unused’ parts of the ‘core’ metaphor, or draw out extended metaphorical relations, they still share a common set of principles which will ultimately be shown to be unsuitable for our purposes. If we want to find a new way of thinking about MEMORY, a good way to do this will be to look for new central metaphors that do not rely on the ‘standard’ construction.

Roediger and Memory Metaphors:

Speaking specifically about the use of metaphor in memory, (Roediger, 1980) discusses how a spatial metaphor, which is generally popular in natural language, had been carried over into studies in cognitive psychology of the time. Roediger explains that “we speak of *storing* memories, of *searching* for and *locating* them. We *organize* our thoughts; we *look for* memories that have been *lost*, and if we are fortunate we *find* them. The term *recollection* implies that memories are objects left at some earlier time that are now collected again” (ibid, 232, original emphasis). He suggests that this creates two main assumptions for memory studies: first, the memories are thought to be discrete

objects found in particular locations in the mind space and, second, that one must search for and find these memories in order to recall information.

After an extensive presentation and discussion of psychological understandings of memory in the century preceding his writing, Roediger demonstrates that a substantial majority of views take on one or both of the assumptions described above: that is, an assumption of spatiality and an assumption of the necessity for search. He also notes the historical progression of these central assumptions in line with technological developments and discoveries over time (e.g. wax tablet to libraries, to dictionaries, to gramophone, to tape recorders, to switchboards, to keysort cards, computers and holograms)⁴⁶. However, he notes (similarly to L&J), that the continued reliance on these predominant assumptions reinforces study of the same aspects of our understanding and may be obscuring new modes of investigation and new findings. As such, he suggests that it would be worthwhile to explore some alternative theories that would allow for new experiments and understandings and proceeds to explore some of the more prominent ones of this time.

Ultimately, he categorizes these alternative theories into two main groups: those which utilize analogies that don't rely on spatiality or search, and those that don't rely on any kind of concrete analogy. The latter types of theories are thought to be those that rely on verbal or mathematical theories that provide abstract descriptions of memory. As Roediger does not dwell on these options, I will not either and will, instead, focus on

⁴⁶ We could add to this memory as writing, as book, Bologna luminescent stone, engraving, phonograph, camera obscura, daguerreotype, photographic plate or photograph, and more (see Draaisma, 2000)

Roediger's findings in the case of the former sort of alternative theories. For these, he ultimately demonstrates that they, too, may still fall into the notions of spatiality and search, but through different means.

Though Roediger discusses a wide swath of alternative metaphors, I will focus on three main groups. The first is (Craik & Lockhart, 1972) levels-of-processing framework, where "memory could be considered a by-product of perceptual processing". Their view suggests that perception of a stimulus "proceed[s] through a series of analyses or processing stages, from early sensory processing to later semantic and associate operatives" (Roediger, 1980, 240). Though the metaphor of levels of processing is not meant to be understood in search-terms, Roediger points out that the levels are conceived as serial and spatially related to each other. As such, it could still be understood as a version of the predominant model.

Another important alternative analogy is that of auditory signal detection, whereby the memories are retrieved not via search but by matching 'resonances' between informational cues and desired memories. An important example of this is (Lockhart et al., 1976) who likened the process of retrieval to a tuning fork: in the same way that when a particular frequency is sounded in the presence of several tuning forks, the fork with the same frequency will sympathetically resonate, Lockhart et al. suggest that a given memory might 'sympathetically resonate' in the presence of the right cue. Here too, this theory doesn't rely on the notion of search for recall. However, it does seem to rely on some sense of spatiality, or memories contained within a given 'mind' space that can still be 'in reach' of each other.

However, a key alternative example that Roediger brings up is (Bartlett, 1995) (which we will come back to in much greater detail in several of the coming chapters). In short, as Roediger puts it, Bartlett's "basic idea is that remembering involves a construction of memories from available information, rather than a verbatim reproduction of the contents of memory" (Roediger, 1980, 239). Indeed, Roediger understands Bartlett's view to be 'deliberately abstract', at least partially so as not to imply a spatial metaphor. However, according to Roediger, Bartlett's ideas were then picked up by others, such as (U. Neisser, 1967), who then introduced new analogies for construction to develop the view. In particular, Neisser compared the process of remembering to a paleontologist reconstructing a dinosaur. He suggested that, in the same way a paleontologist can construct a complete model of a dinosaur from a few bones (and some knowledge of how they should be put together), we can also reconstruct memories from a few details and/or a general impression. Here, again, we can see that, though the notion of search is not used explicitly, the construction of and relation between the bones can be understood as at least partially a spatial relation. Furthermore, it seems reasonable to think that the bones would have to be stored and retrieved somewhere in order for (re)construction to happen.

So, even in these alternative cases, Roediger's examination highlights the long-standing dominance of the spatial and search conditions in memory. Though Roediger's notions of spatiality and search are important and useful in demonstrating the various differentiations amongst the views he describes, in the next chapters I ultimately consolidate these various views under the umbrella of a more contemporary parlance,

namely a Storage and Retrieval (S&R) model of memory. Roediger's 'spatiality' relates to the notion of 'storage' in so far as, under the S&R model, memories must be, in some sense, kept and maintained in a way that usually has some spatial implications: be it like boxes in a warehouse, or files in cabinets, documents in files, or what have you.

Similarly, Roediger's 'search' is related to 'retrieval' though not identical. Within S&R parlance, memory must be 'searched' in order to identify and then 'retrieve' the appropriate memory. In this sense, I believe the 'search' assumption implies a 'retrieval' function as well and is not so far off from what I think Roediger intends.

In the following chapter, I will examine a strict S&R view and lay out its deficiencies for explaining procedural memory and its failure to account for empirical findings, such as Bartlett's, that cast doubt on the general efficacy of the S&R metaphor. The (re)construction metaphor, on the other hand, does highlight some new aspects of memory in general and its application to procedural memory is also fruitful. However, as Roediger already hinted towards, a similar fate lies in store for (re)constructionist views, in so far as they still maintain some fundamental S&R assumptions, especially relative to 'storage'. Even so, there are more contemporary versions of (re)constructionist views that introduce some ways around these worries that are, as yet, incomplete, but useful for our purposes. Ultimately, we should end up with some promising starting points for a how a new set of metaphors and new understanding of procedural memory.

Chapter 4: **Storage and Retrieval Model of Memory**

4.0 - Introduction:

As was stated in the introduction to Part 2, this part of the project is aimed at examining dominant understandings of memory and determining whether and to what extent they apply in the case of procedural memory, or memory of skill. This chapter is aimed at examining one such class of views which, until relatively recently, were dominant in both philosophical and psychological spheres of research: what I am calling the Storage and Retrieval (S&R) model of memory. This model has a huge intuitive appeal and has been a culturally dominant model for (arguably) millennia.

I will begin by laying some of the historical grounding of the S&R model and culminate with a dominant late-20th century iteration of it based in information-processing. From there, I will explain the main subdivisions of memory and demonstrate how S&R is built around, and best suited for explaining, one particular subdivision of memory: semantic memory. However, I will then offer a potential reading of procedural memory under and S&R account and demonstrate its insufficiency via two methods. Firstly, I offer a conceptual analysis of the application of S&R to procedural memory and the ways in which it falls short of describing skill according to our previous criteria. Along the same lines, this prong of the demonstration will also reveal the manner in which the S&R metaphor highlights particular aspects of memory and obscures other key aspects that are important to procedural memory. Secondly, I examine how psychological research and experiments of the last century have proven a strict S&R model to be obsolete, even in its ideal setting of semantic memory and suggest that this

should remove the S&R model from consideration as an understanding of procedural memory altogether.

4.1 – The Storage and Retrieval Model:

4.1.1 - Historical Accounts of Memory:

The name ‘Storage and Retrieval’ is meant to capture what are understood to be the two basic functions of memory: the means of recording and maintaining impressions of past experience in the mind (i.e. storage), and the means of recovering those impressions (intact) for later use (i.e. retrieval). These basic functions have been described in various ways throughout philosophical history, but they have remained largely unchanged throughout that time.

This basic format can be seen as early as Ancient Greek conceptions of memory.

Plato describes memory in Theaetetus, through Socrates, as a ‘block of wax’. He says,

... when we wish to remember anything which we have seen, or heard, or thought in our own minds, we hold the wax to the perceptions and thoughts, and in that material receive the impression of them as from the *seal of a ring*; and that *we remember and know what is imprinted as long as the image lasts*; but when the image is effaced, or cannot be taken, then we forget and do not know. (emphasis added, Plato, 1956)

In other words, Plato’s description suggests that our experiences and perceptions are stamped onto the block of wax, as if from ‘the seal of a ring’ and stored in our minds as long as the imprint on the wax remains intact. In this sense, they are ‘stored’ in the block of wax in our minds. In order to recall or retrieve this imprinted information from storage, one need only pass one’s thoughts over the block of wax and, in a sense, read the memories off of it⁴⁷.

⁴⁷ It is worth noting here that Plato ends up rejecting this notion as an acceptable model of memory. Later on, he introduces the notion of memory as an aviary. The aviary contains a great many birds that fly about

In very similar fashion, Augustine describes his understanding of memory at the very end of the fourth century. He describes memory as a ‘storehouse’, writing,

And I enter the fields and spacious halls of memory, where are *stored as treasures the countless images that have been brought into them* from all manner of things by the senses. *There, in the memory, is likewise stored what we cogitate*, either by enlarging or reducing our perceptions, or by altering one way or another those things which the senses have made contact with; and everything else that has been entrusted to it and stored up in it, which oblivion has not yet swallowed up and buried. When I go into this *storehouse*, I *ask that what I want should be brought forth*. (...) *The vast cave of memory, with its numerous and mysterious recesses, receives all these things and stores them up, to be recalled and brought forth when required.* (emphasis added, Augustine, 1955, Book X, Chapter VIII)

Here, the language of ‘storage and retrieval’ is quite explicit. Augustine quite clearly states that thoughts and experiences are ‘stored’ in memory as if they were items or goods or property in a ‘storehouse’. So long as the items are properly maintained or undamaged in the storehouse, they can be said to be part of one’s memory. When one wants to retrieve or remember something, one need only, as it were, walk into the storehouse and retrieve whatever it is that one is after. Or perhaps, since Augustine writes that he ‘asks that what [he] wants should be brought forth’, one needs to ask to the operator of the storehouse to bring forth whatever it is one is after. In either case, the two

it freely. The birds are meant to represent memories contained within a regulated space, but not as neatly pinned down and stored as on the block of wax. Still, this metaphor could also be understood as a version of the storage and retrieval model – information (birds) is stored in the aviary and must be retrieved (caught) in order to remember. The aviary, however, requires a much more active sense of retrieval than the case of the wax tablet, as one must identify and catch the unruly bird, and then let it go back into a rather haphazardly organized cage once one is finished with it. By the end of the dialogue, however, Plato rejects this aviary model, as well, and seems to settle on a notion of knowledge as constructed by smaller bits of component knowledge – where knowledge is constructed like a cart is from its parts, or like words are constructed from letters. In this way, then, memory might be a process of finding or gathering all of its requisite parts and putting them together in the right way. This latter view starts to look more like the view that will be discussed in the next chapter (i.e. reconstruction). However, it is not the aim of this chapter to settle an interpretation of Platonic notions of memory – only to show that understandings of this sort have been considered and discussed since at least Plato’s time and that they have remained influential.

main mechanisms of memory are understood to be some manner of ‘storage and retrieval’.

Even as far on as the late seventeenth century, John Locke describes memory as a ‘storehouse’, even maintaining the same turns of phrase used over a millennium before.

He writes,

This is memory, which is as it were the *storehouse* of our ideas. For, the narrow mind of man not being capable of having many ideas under view and consideration at once, it was necessary to have a *repository*, to lay up those ideas which, at another time, it might have use of. But, our ideas being nothing but actual perceptions in the mind, which cease to be anything when there is no perception of them; this laying up of our ideas in the repository of the memory signifies no more but this,—that the mind has a power in many cases to revive perceptions which it has once had, with this additional perception annexed to them, that it has had them before. And in this sense it is that our ideas are said to be in our memories, when indeed they are actually nowhere;—but only there is an ability in the mind when it will to *revive them again, and as it were paint them anew on itself*, though some with more, some with less difficulty; some more lively, and others more obscurely. And thus it is, by the assistance of this faculty, that we are said to have all those ideas in our understandings which, though we do not actually contemplate, yet we can bring in sight, and make appear again, and be the objects of our thoughts, without the help of those sensible qualities which first imprinted them there. (emphasis added, Locke, 1999, Book 2, Chapter 10)

Here, again, the basic concepts of storage and retrieval can be seen, though they are more refined. The notion of ‘storage’ remains largely the same - ideas and past perceptions are said to be stored in the repository of memory. However, Locke suggests that they are not actually stored in our minds in any literal sense since they “are actually nowhere”⁴⁸. What is ‘real’ is our ability to retrieve, or as he says, ‘revive’ these past perceptions. In other words, there is no true ‘storehouse’ in the mind, only our ability to ‘revive’ past experiences and recognize them as such. It is this faculty which makes it possible for us to say that these ideas or perceptions are in our understanding or

⁴⁸ By saying they are ‘nowhere’, I believe he means to suggest that they cannot be understood as occupying a particular space in the mind and must, in some sense, be distributed throughout it.

memories. Though Locke focuses on the faculty of ‘retrieval’ and seems to suggest that this is the defining feature of memory, he does still utilize the imagery of memory as ‘storage’ and an understanding of whatever can be said to be in our memories as ‘stored’ there.

This brief historical survey illustrates that, though there were various iterations and refinements of the theory over time, the understanding of memory as a process of ‘storage and retrieval’ has been a common one over time and has proven to be enduring. This pattern continued into the 20th century. However, arguably the most substantial changes and revisions of this model came during this time, as the model was adopted into experimental psychology and largely translated into the language of information processing. During this time, it came to be part of the developments in artificial intelligence and computing and, as such, began working its way back into the philosophical understanding in this updated form.

4.1.2 - a Late 20th Century Understanding of Storage and Retrieval:

In what follows, I utilize descriptions of the S&R model as it was understood in 1970 by Endel Tulving and Stephen Madigan. Tulving, in particular, is regarded as a central figure in psychological memory studies of the 20th century. He subscribed to the S&R model in his own work and he and Madigan described it as the basis of one of two main camps in memory study at the time (the other being the study of verbal learning) (Tulving & Madigan, 1970). I find their description of S&R to be particularly useful for several reasons: it is concise, comprehensive, and representative of a great many iterations of this model both prior to and after their own writing. Given that the

description in their paper has to do primarily with experiments and studies relating to verbal processes, I will adapt their language slightly to make it more generalized.

As Tulving and Madigan (henceforth T&M) explain, much of the 20th century version of the S&R model comes from the contemporary developments around computers and the early stages of research into artificial intelligence. As such, much of the language used in the model is that of information processing. The S&R model understands the process of memory to have four core components: coding, storage, cueing and retrieval. Coding and storage are closely related, as are cueing and retrieval. T&M explain the basic set-up as follows: “When a subject sees an object or a picture, or hears an isolated item or a sentence, appropriate information is entered into his memory store. (...) Stored material can be ‘utilized’ by the subject in recall – or recognition tasks. Utilization of stored information, or retrieval, is often likened to a search and decision process” (ibid, 440). On this model, then, having a memory of something involves *both* having some version of the perceived information in one’s memory store and the successful retrieval of said information later on.

Let’s begin with coding. In order for perceptual information to be stored in memory, it must undergo some manner of translation into what could be described as an appropriate memory ‘format’. This might involve the creation of corresponding representations, associations, or facsimiles of information being created and stored but, in any case, there is a difference between the physical stimulus presented to a subject and the information stored in memory. T&M detail two types of coding. The first, substitution, involves the “replacement of the input stimulus by another symbol, the code,

together with a general decoding rule” (ibid, 462). In other words, information is written into memory as a ‘code’ that can be decoded with a general rule already present in the system when the time comes. The second, elaboration, involves the “storage of additional nonredundant information with the [information] unit to be remembered” (ibid, 462). In this case, a unit of information is encoded along with some extra information that is relevant to the core unit. This means that previously encoded units of information can be ‘edited’ to include new information. This process of encoding is a key factor in the process of memory in so far as it determines what will be available for recall later⁴⁹.

Storage, however, is distinguishable from encoding, though very closely related. After information gets encoded, it must be kept and maintained so that it might be used later. T&M have little to say about the mechanisms by which this kind of maintenance is done, as, indeed, does little of the literature on this subject. Instead, what is primarily discussed are the different types of storage, namely short-term and long-term storage. In short, it is thought that all information presented to a subject first enters the short-term store and some of it, depending on several factors (including available capacity at a given time as well as rehearsal), will be transferred to long-term storage. This means that whatever information enters short-term memory and is not transferred to long-term storage is, in some sense, lost.

⁴⁹ Newer versions of S&R (such as (Schacter et al., 1998) also add that encoding should be able to, first, bind together all relevant information to a given episode or event, and second demarcate the limits of individual events or episodes such that they don’t bleed into each other. In other words, demarcating separation between various ‘bundles’ of information is a key aspect of encoding.

As a partial explanation of this, T&M cite the work of George Miller and his 'leaky bucket' hypothesis. They explain that,

(...) as items enter the short-term store, a point is reached at which old items 'leak out' as fast as new ones are put in, or existing memory traces fade away as fast as new ones are created. (...) The short-term store is assumed to have a limited capacity for holding information, and once that capacity is reached, further incoming items displace those currently in the short-term store according to some probabilistic schedule (ibid, 469).

In other words, short term memory can only hold so much information at a given time. If its resources are not in great demand, it can process more information, thereby increasing the chances of said information being transferred into long-term storage. However, if its resources are in full use, it is very likely that information will only be held for a certain amount of time before being 'let go' in favor of newer information. In this latter case, some information that comes into the short-term store will not be transferred to long-term store.

Another important factor in the transfer from short- to long-term store is 'rehearsal': that is, use, especially repeated use, of information. The idea here is that if a piece of information is used either as a central component in the performance of a task, or repeated sufficiently, it has a strong likelihood of being transferred to long-term memory. T&M explain that, at the time they were writing, the effects of rehearsal were well known and reliably reproduced. However, there was not much known about the characteristics or properties of the process (ibid, 469). This largely remains true today.

Once something is properly encoded and stored, the process of retrieval can begin. On the S&M model, all retrieval begins with a cue. In other words, the system must be triggered in the appropriate way to begin a search of its storage space to find relevant pieces of information for a given task. T&M explain that there are many types of

effective cues and that some are more effective than others. The effectiveness depends on some widely agreed upon factors. First, it depends on the nature of the encoding process for a given piece of information. That is, the more effectively something has been encoded, the more effectively it can be cued, or be used as a means of cueing other related pieces of information. A second factor is the extent to which a unit of information is connected or associated with other unit(s). The stronger the connection, the more effective the cue, and vice versa. A third factor is the number of items connected with a given cue. The more items associated with a cue, the less effective it will be. If there are only a few associations, the system can pull them all up with ease when cued. But larger batches of associations will require more effort to cue and will be cued less effectively. Finally, the strength of preexisting associations between a cue and a unit of information will impact its effectiveness. Preexisting associations tend to be stronger and more effective than newer ones (ibid, 465-466).

Finally, retrieval, according to T&M, has two fundamental principles. The first is that, in order for something to be retrieved, it must be present in storage. The second is that failure of retrieval does not necessarily imply that the relevant information was not in storage. In general, they claim that much of the literature treats retrieval as a set of two subprocesses, namely search and decision. When a system starts a search, or scanning processes, it finds “access to stored information of all kinds” (ibid, 461). That is, the system’s memory stores are scanned in search of any and all relevant pieces of information. Decision, on the other hand, “(...) compares some property of the ‘found’ information such as its ‘familiarity’ or its ‘list tag’ with a hypothetical entity called the

‘criterion’” (ibid, 461). This means that once the system has identified a piece of information as relevant, it will compare it with the search criteria to determine whether it meets them or not. T&M go on to explain that, “[d]epending on the results of such a comparison, the information is used to produce a corresponding response or is rejected, whereupon the search process continues until the desired information is found and identified as such, or when it grinds to a halt for as yet unknown reasons” (ibid, 461). In other words, if the relevant unit of information meets the search criteria, it will be retrieved and used in the performance of a task. If it does not meet the criteria, it will be rejected and the search will begin again until a unit is found that matches the search criteria or, for whatever reason, the system stops the search.

This constitutes a brief overview of the general structure of the S&R model. T&M, as well as Tulving himself in other works, cite several historical philosophical figures to lay the intuitive groundwork of basic conceptions of memory in order to develop their own theories (Tulving, 1983; Tulving & Madigan, 1970; see also Roediger (1980) who also draws a direct link between S&R type accounts and philosophers as far back as Plato). In this way, they started from historical philosophical explorations that relied on a basic version of a S&M, and developed a new, updated, 20th century version of it. Within the philosophical literature, some version of S&R, often the information-based version, has been widely adopted and gone largely unchallenged (see empiricist theories of memory (Byrne, 2010); epistemic theories of memory (Adams, 2011; Audi, 2002; Williamson, 2000); causal theories of memory (Bernecker, 2008; Debus, 2017; Michaelian, 2011; O’Brien, 1991; S. Robins, 2016a; S. K. Robins, 2017) . In addition to

its wide acceptance, it has also been widely thought to describe the most basic mechanisms of memory and so to apply to all cases and types of memory across the board.

4.2 - Semantic vs Procedural Memory:

Having now walked through the basic tenet of S&R, we can begin to see how this might be applied to different kinds of memory. As such, a basic primer on the types of memory to be discussed in this paper is in order. Both Philosophy and Psychology have coined many terms to denote the same or similar types of memory. However, the standard taxonomy of memory, developed in historical conjunction between the two disciplines, separates between two main classes of memories (Michaelian & Sutton, 2017). The first is what philosophers call ‘declarative’ memory, while psychologists call it ‘propositional’ memory. This larger grouping contains two further types of memory, semantic and episodic. Semantic memory has to do with remembering facts, conceptual information, or propositions more generally (i.e. Riverside is the biggest city in the Inland Empire), while episodic memory has to do with personal, direct experience of events (i.e. I remember moving to Riverside). For the purposes of this paper, I will not discuss episodic memory and will focus, instead, on semantic memory.

The second main class of memories is what philosophers call ‘non-declarative’ memory, while psychologists call it ‘procedural’ memory. This kind of memory is typically understood to relate to memories of acquired skills and habit memory (i.e. I know how to play an instrument). The notion of habit, I imagine, plays a large role in the origin of the name ‘procedural’: these kinds of memories and skills are thought to be so

well rehearsed and so well engrained that the subject need only run through an established 'procedure' to perform the task, as opposed to having to think consciously about the steps required or bringing their attention directly to the action.

Endel Tulving, a very prominent psychologist of the second half of the 20th century, provides a neat and tidy characterization of the differences between propositional and procedural memory.

- (1) procedural memory is not thought to be symbolically expressible, unlike semantic and episodic memory;
- (2) procedural memory is demonstrated through the performance of the skill in question, while propositional memory is demonstrable through a wide array of behavior;
- (3) procedural memory does not involve veridicality while propositional memory can be talked about in terms of truth or falsity;
- (4) procedural memory is acquired through extensive practice while propositional memory can be acquired in a one-off manner;
- (5) and, finally, procedural memory involves an absence of thought and relies on a sort of automation while propositional memory usually requires directed attention.

(Tulving, 1983)

These points are not universally agreed upon. For instance, some might think that procedural memory could, in fact, be acquired in a one-off manner or that both propositional and procedural memory could be understood as 'thoughtless'. However,

Tulving's picture is a common one and at least serves the purpose of painting a picture of how procedural memory is usually understood.

4.3 – Understanding Procedural Memory in terms of S&R

S&R seems especially well suited to explaining semantic memory – the notion that we remember facts, propositions, and conceptual information through a process of coding, storage, cueing and retrieval seems relatively intuitive. In the same way that, when working on a computer, one must create, save, find and open a file with all your relevant information, the same could be thought to be happening in our memories with facts and concepts. Suppose the piece of semantic information in question is that Mozart's birthday is January 27, 1756. For it to be said that I remember this would involve, first, it having been properly encoded and stored in my memory system; that is, it having been 'converted' into the appropriate 'file type' and then being stored in memory without sustaining damage, being lost track of, or being deleted. Then, when appropriately cued, I should be able to retrieve that information intact and be able to tell you that correct date. This kind of semantic information fits the model very well and probably seems quite intuitive.

In terms of procedural memory, then, it seems that an S&R account would suggest that it must work along the same four basic parts: that is, there must be some sense in which it is coded and stored in memory, and then appropriately cued and retrieved when necessary. For instance, knowing how to play an instrument could be understood as a collection of semantic information put together in the right way. Perhaps semantic information such as 'holding my fingers in *this* way creates *this* note' or

‘playing *this* way will create *this* tone’, etc. Or similarly, a procedure like walking could be understood as a series of semantic facts such that the subject might think, “raise foot to *this* height; push it forward; balance body weight at *this* point between feet; put foot down; push weight *this* far forward; repeat with the other leg, etc.”. If all of this is true, then remembering a skill could still be understood to rely on encoding, storage, cueing and retrieval (i.e. S&R).

This is perhaps not very dissimilar from understandings of skill that we covered in previous chapters, especially those discussed in epistemological understandings of skill. For example, in so far as knowledge-that lines up with semantic memory, and knowledge-how lines up with procedural memory, it might not be unreasonable to guess that (Stanley & Williamson, 2001) would endorse the idea that procedural memory ultimately reduces down to a kind of semantic memory. That is, that procedural memory ultimately comes down to storing and retrieving semantic information about the skill in question. Furthermore, taking into account the more up-to-date versions of their thinking, we could consider (Stanley & Krakauer, 2013), who suggest that skill crucially involves knowing a proposition of how to initiate an action, and (Stanley & Williamson, 2017) who suggest that skill involves a disposition to know how to guide the action in question (itself understood to involve semantic or propositional information). Though both of these later papers highlight that there are other aspects of skill, namely perceptual and motor acuity, that are not captured by their suggestions, they suggest that these other aspects are not relevant to issues of knowledge. Taken together, these versions, too, might suggest that procedural memory ultimately involves storing and retrieving

semantic information throughout performance. If the suggestion that procedural memory reduces to semantic memory suggested by the work of Stanley, Krakauer and Williamson is true, it seems plausible that S&R would be in a good position to explain procedural memory.

However, the potential equivalence starts to break down as we examine an application of S&R to cases of procedural memory. As we will see, S&R could be understood to capture the semantic elements of memory of skill but tells us very little about elements of our previously established RRRR model and elements of perceptual and motor acuity. Insofar as procedural memory, then, deviates from semantic or episodic memory, S&R will prove to be insufficient for our purposes. Furthermore, we can see that the S&R model is predisposed to hide some of these other elements or force an understanding of them that would be compatible with the model. To help demonstrate this, let us consider the following question: what would it take for someone to remember a skill?

4.4 – Why S&R isn't equipped to explain Procedural Memory:

Remember our question: what would it take for someone to remember a skill? On a strict S&R account, having stored information about the skill in question is obviously important, but not sufficient for memory of a skill. It seems important that one be able to *perform* the skill successfully to be said to remember how to do it (recall, also, that this is one of Tulving's characteristics of procedural memory). For instance, suppose that Hannah used to be a very active cyclist but hasn't ridden her bike in many years. Even so, Hannah claims to remember how to ride her bike. If, when she gets on the bike, Hannah

rides her bike successfully, one would likely think she was right – she did remember how to ride her bike. However, if she is ultimately unable to successfully do so, this would suggest that she did not actually remember how to do it⁵⁰. She may remember *something* about riding a bike – some smaller aspect or fact or experience - but being unable to perform successfully suggests that whatever she remembers is incomplete. This could be the case for any number of reasons: perhaps said information wasn't stored properly in the first place; maybe some of that information has been lost from her storage over time; maybe it was stored and maintained properly but wasn't retrieved properly. Regardless of which of these is the case, it seems important that the process of both storage *and* retrieval go right in order for her to be said to remember the skill, with retrieval being key.

This being the case, an S&R proponent might offer the following answer to our question: one would have to retrieve a set of actions, or perhaps a procedure, from storage and implement them as faithfully as possible. If one could do that, one could be said to remember a skill. To flesh this out, let's consider a few examples around professional soccer player, Megan Rapinoe. Suppose Rapinoe practices penalty shots at the end of every one of her training sessions. Let's assume that, today, she is practicing an established skill such as bending the ball into the top-left corner from the penalty spot. On a strict S&R reading, what it would take for Rapinoe to remember the skill of bending

⁵⁰ Here we are talking about normal circumstances for Hannah – that is, we are assuming that in the intervening time her body hasn't changed significantly (for instance, she hasn't gained or lost a large amount of weight, or became paralyzed, or an amputee or what have you). We are also assuming that there are no external factors (like wind or someone pushing her) that would prevent her from being able to ride her bike normally. As such, we are assuming that conditions are as they usually are for her and there are no circumstantial hinderances to her performance.

the ball would be to retrieve all the relevant information from her memory storage. This might include a procedure such as starting at the ball and taking x number of steps back at a stride length of y ; running back up to the ball at a pace z ; hitting the ball with the precise part of her foot, f , at precisely the right point, p , with precisely s amount of swing. If all of this could be accurately retrieved, Rapinoe would successfully remember her skill and would be able to bend the ball.

However, even in this repeated, regulated scenario, it seems unlikely that retrieving this procedure would be enough to encapsulate Rapinoe's skill of bending. Each time she kicks the ball will be a little bit different, even if the general procedure remains the same. For instance, depending on her practice prior to the penalty shots (either that day or in previous days), she may be tired or sore in different parts of her body during different sessions; similarly, with every kick she makes during this particular session, she may get a bit more tired, a bit more depleted. As the dynamics of her body shift from session to session, or within a given session, she will have to adjust and regulate herself to ensure that she is in line with the procedure and able to perform it successfully. That is, if she is to maintain the *resilience* and *reliability* of her skill, she will need to be *responsive* to the dynamics of her body at any given time. As such, if she must incorporate all of this into the established procedure to make it viable in this instance, this might mean that there are other aspects of Rapinoe's skill of bending that are not encapsulated in the retrieved procedure.

These worries can be significantly heightened as we move away from controlled environments, like practice penalty shots, to more realistic scenarios of play. Suppose

Rapinoe is taking a bent shot on goal in the middle of a competitive game. She is no longer taking the shot from the usual penalty spot, is moving rapidly and dodging opponents to get to a ball that her teammate is passing to her at high speed. If Rapinoe retrieves her established procedure for bending the ball, we can see that it will necessarily require adjustments which will depend on where she is on the pitch, how her body is situated as she approaches the ball, where the opposing players are, and where the goalkeeper is. Each kick she could take from each plausible point on the pitch would necessarily be different from those practiced in her established procedure – she could not implement the procedure identically from each one of these spots or even each time she kicks the ball. Here, too, if she is to maintain the *resilience* and *reliability* of her skill, she will need to be *responsive* to the dynamics of her body and everything around her at any given time.

In both these cases, Rapinoe would need to supplement her established procedure with something in order to make it responsive. A likely solution here would be for her to draw on past experience or information from memory in order fill out her procedure and be able to successfully perform her bent shots. That is, she might need to recall general rules or set up from training, similar instances of play that would give her relevant guidance or use those memories to put together a new plan altogether. For instance, she may recognize that her right leg is fatigued and cannot produce the usual amount of power. As such, she will need to find a way to compensate for it. She might then remember that she has done this successfully in a number of practice sessions by incorporating a greater twist of her hips into her swing that would match her usual

kicking strength. Or similarly, in the case of real game play, she might recognize that her current positioning, as well as that of her opponents, is similar to another instance of play where she successfully bent the ball and applies those findings to her current case. In this way, given her diagnosis of the state of her body and of play, she can draw on her past experience to inform how she will modify and execute the original procedure. On an S&R model, then, Rapinoe would need to retrieve any and all relevant instances of practice or play and incorporate her findings into the original procedure for bending.

Taking stock, so far this suggests that, on an S&R account, Rapinoe's skill at bending would need to include more than just the original procedure: in addition, it should also be able to incorporate relevant stored information about how to modify the original procedure based on other past experience. Indeed, this seems to be what Stanley, Krakauer and Williamson are suggesting in their aggregate view: knowing how to start and being disposed to utilize knowledge to guide the rest of the action.

However, this raises a problem discussed in a previous chapter, namely a problem of association. (H. L. Dreyfus, 2002), drawing on Merleau-Ponty, mentions that there are many dimensions along which experiences can be similar to each other and it's not clear why a given memory would be called up as the appropriate one. For instance, if Rapinoe searches her memory for a similar instance of play, it seems unlikely that she will find one that is 100% similar to her current situation – no matter how similar, no two instances will ever be exactly identical. Instead, she may retrieve from memory several different memories that draw on different aspects of her current circumstances but are all equally relevant or related: maybe she draws up 4 different memories, each is 50%

relevant but relevant to different parts of the experience. It's not clear how she might differentiate between them and decide which one to use in her present circumstances⁵¹. Should she combine them all? Which parts? Only those that overlap? Maybe some of them are actually irrelevant or unhelpful to the circumstance, though they were 50% relevant – does she include this in the amalgam?

More importantly, it seems plausible that no one actually remembers all the relevant instances of practice or performance that contributes to their skill – neither toward a core procedure nor other relevant instances of performance. Indeed, it seems more likely that, the longer one is a practitioner of a given skill, the more likely one is to forget large swaths of practice or play at the beginning of one's skill acquisition. Someone like Rapinoe, who has been playing soccer professionally for more than twenty years (and who likely started learning much earlier than that), is unlikely to remember every instance of practice in the early parts of her career, especially those that laid the foundation of her skill. Further, it seems unlikely that she remembers every instance of practice or play over the span of her two-decade long professional career. Even so, it seems that this does not impact her performance in any significant way, and she is still able to adjust her performance in accordance with her circumstances of play. This then means that being able to retrieve relevant instances of past experience is not crucial for skilled performance which, in turn, casts doubt on the strict S&R understanding of skill described above. If remembering a skill requires the retrieval of a set, or static, procedure

⁵¹ Indeed, this might be understood as a version of the Frame Problem applied to cases of memory rather than beliefs.

from memory *and* being able to retrieve other relevant experiences with which to modify it, it seems unlikely that one would be able to perform at a high level.

An S&R proponent might respond by suggesting that one doesn't actually need to store and retrieve all of one's relevant past experiences over time, only key ones. Indeed, maybe these could be somehow tied to the central procedure and then recalled along with it. However, this may raise similar worries to the those described above in terms of the problems of association: how would those memories which are most important be identified? One potential problem, here, is that the S&R proponent's line seems to suggest that there are some key, singular moments of insight or development that happen in semi-linear fashion that could be isolated from others. That is, it relies on the idea that, as one acquires skill, one has moments of acquisition that are particularly important for developing one's skill, that those moments are completely distinct from all those that come before it, and that, as one gets better, one builds on the important acquisition in order to get to the next one. However, as we saw in earlier chapters (especially Chp. 3 on walking), it seems much more likely that the acquisition of skill is not linear in this way and that there would be too few of these shining moments of development that could constitute a skill at large. That is, that skills are not acquired in a 'step-like' fashion and that it is not clear that one moment of insight or development can be clearly differentiated from another.

Ultimately, it seems unclear how all of this accounts for several elements of our RRRR criterion and how perceptual and motor acuity fit in to skill. Recall that the RRRR (resilience, reliability, responsiveness and robustness) account of skill aimed to identify

and explicate several axes along which skilled performance could be characterized. To begin with, given everything we've seen above, an S&R account does not seem to offer a satisfying understanding of how resilience, reliability or responsiveness can happen, since it does not seem to offer an explanation of how experience can be accrued over time and appropriately deployed to be able to perform skillfully across these various dimensions. Indeed, the axis that seems furthest from reach is that of responsiveness: given that it is unclear how past experience can come to bear, the ability to take one's body and surroundings into account and being able to respond appropriately in light of multiple considerations seems difficult to achieve.

Furthermore, implicit in much of the discussion of RRRR was an understanding that perceptual and motor acuity are part of what allows for high marks along each of these axes. Though Stanley, Krakauer, and Williamson do allow that these are relevant aspects of skill, they deny that these are relevant to discussions of knowledge and, presumably, memory. However, it seems that there must be some sense in which they are remembered. If memory is understood very broadly as the influence of past experience on present and future behavior, it seems there must be some sense in which changes in perceptual and motor acuity garnered through practice and past experience could be understood as being remembered. However, the S&R model does not seem to allow for this at all.

All of this suggests a deficiency of the S&R model that I think is at least partially attributable to the structure of the metaphor. Recall the work of (Lakoff & Johnson, 1980) in the introduction to this section, who suggest that whatever grounding metaphors we

use will both naturally highlight some aspects of the phenomenon they are describing, and hide others. In general, I suggest that the S&R model highlights the role of propositional and declarative knowledge in memory, while obscuring the possible contribution of non-declarative and procedural knowledge – this is at least part of why the view fails for our purposes. Indeed, the only way it can accommodate the influence of the latter is to understand them in terms of the former. However, in trying to apply the S&R model to memory of skill, we can see the impact of this structural bias in its inability to account for several key aspects of the target phenomenon. In this way, the metaphor of S&R is ill-suited to our purposes of explaining procedural memory (or memory of skill).

4.5 – Dissent within Psychology

Beyond being inappropriate for our purposes (relative to explaining procedural memory) there are reasons to believe that the S&R model is simply empirically false. Even as early as the 1930's, there has been dissent within psychology around an S&R model of memory, in general. For instance, F. C. Bartlett published the results of several experiments in 1932 which cast doubt on S&R, understood as a purely 'reproductive' understanding of memory, and laid the foundation for what are called (re)constructive models of memory (Bartlett, 1995). Indeed, the foundations laid by Bartlett's work set the stage for, and were later substantiated by, more contemporary research. This ultimately led to a nearly discipline-wide rejection of a strict S&R model in psychological memory research. I will not, here, dwell on more contemporary research as it will be more fully developed in the next chapter (dedicated specifically to reconstructive views). Instead, in

what follows, I'll explain two key experiments Bartlett performed that helped demonstrate the insufficiency of a reproductive S&R model which I think are sufficient to prove the point.

S&R is usually understood to a model that requires accurate reproduction of memories: they must be stored, maintained, cued and retrieved in such a way as to maintain their integrity and reproduce all the information that was originally encoded. In general, Bartlett's results across both experiments showed several important things: first and foremost, both suggest that accurate remembering, as suggested in an S&R model, is the exception, not the rule, and that even small amounts of simple material is rapidly affected by transformation. That is, even with relatively simple material and after a short span of time, subjects' memories will, more likely than not, be at least slightly distorted. In both cases, accuracy decreased and transformation increased the further out from the initial sitting the subject was interviewed.

4.5.1 – F. C. Bartlett – Experiment 1 - Faces of Military Men:

Twenty subjects were provided with five cards, each with an image of the head of a military man, faced down on the table, always in the same order. All the heads were similar, but had their own unique, identifying characteristics. Subjects were instructed to look at each card for 10 seconds, trying to remember as much about it as possible, so that they could answer questions about it later. Once the observation period had passed, subjects were occupied in conversation with the experimenter for half an hour, and then asked questions about the cards. They were then asked questions about the same cards one or two weeks later, and then at increasing intervals between visits. Subjects were not

shown the cards again until it was clear that they would not be able to come back for continued experimentation.

Results:

In general, subjects misremembered several features of the faces (their order of presentation, which direction the face was facing, details of their faces, etc.). Bartlett noted that the transformation seen in later memories was informed by one or, mostly likely, both of the following: 1) what kind of method of recall the subject was using (visual or linguistic) and 2) interests or feelings (either those personally held by the subject or those stipulated more widely by a culture or society). To begin with, those who relied on visual imagery for recall (i.e. those who relied on recalling the images as such) were more prone to making mistakes about the order of the presentation of the images and were more likely to import details external to the images. Similarly, they were more likely to express confidence in their recollections (which turned out to be completely divorced from their accuracy).

On the other hand, those who relied on vocalization or language as their primary method of recall (i.e. those who talked through their recollections or reasonings and often inferred parts of their recollection in this way), were more likely to rely on an “*idée directrice*”⁵², or a “guiding idea”. That is, they generated some rule or a guiding principle of the arrangement they saw (i.e. “it seemed to me that all the faces started looking left and by the end transitioned to the right”), which then became the dominant aspect of the memory and shaped the subject’s later recall. Subjects who used this method were more

⁵² This is a notion Bartlett borrows from Binet (ibid, 52)

likely to have an attitude of uncertainty (which was also divorced from the accuracy of their recollections).

4.5.2 - Experiment 2 - Reproduction of a story:

In this case, Bartlett presented twenty subjects with a translation of a North American folk tale, called the War of the Ghosts⁵³. Each subject was allowed to read the story twice at their own pace and was then asked to reproduce it after fifteen minutes. Subjects were also called back to reproduce the story at increasing intervals after the initial meeting as the opportunity presented itself. Subjects came back and attempted to reproduce the story 3 or 4 times over intervals of days, weeks, months, and even years (one subject came back 10 years later).

Results:

Here, too, Bartlett's results suggested that accuracy of reproduction was the exception, not the rule, even at the first reproduction 15 minutes after reading the story. Most subjects altered several features of the story, for example shortening it, reproducing an altogether different style or tone, inventing or including extraneous detail, changing names, changing the order of events, etc. Indeed, subjects often were quick to notice and remark upon things like the choppy style and rhythm of the original story or its strange, supernatural ending, and yet still failed to faithfully reproduce them even shortly afterwards.

⁵³ From what I can gather, it is an indigenous folk story from the Pacific North-West of the United States. The story mentions a Kalama River, and there is a Kalama River in what is now Washington State.

Over time, those who were frequently called on to reproduce the story commonly settled into a general form or outline of their version of the story, or what Bartlett calls a ‘stereotype’ of the story. These stayed largely the same over repeated reproductions. However, those who only infrequently reproduced the story had a contrary tendency: they were more prone to omitting detail, simplifying the story, and otherwise transforming details. This process of transformation seemed to have been able to go on indefinitely, as long as the subjects were unaided or corrected by the researcher.

However, long-distance remembering (or remembering after long periods of time) seemed, in some instances, to increase elaboration of the story, importing extraneous detail into it, and/or inventing new details. Long-distance remembering seemed to happen most often in one of two ways: on the one hand, in a way similar to the “*idée directrice*” mentioned above, subjects maintained a general setting of the story which was expressed through their attitude towards it. On the basis of this general setting, they were able to construct and/or infer a great deal of detail. However, this detail usually had little to do with the original story and more often reflected the subject’s personal interests and tendencies in their normal life. There was some indication that the influence of a subject’s attitude grew stronger as time went on. On the other hand, some subjects only remembered one or two details that they found striking about the story and little, if anything, else.

Takeaway:

In both of these experiments, it seems clear that subjects were not engaging in a strict S&R manner, which would suggest that the subject need simply find the relevant bit

of information through appropriate cues, etc. and reproduce it as accurately as possible. In both the experiments described above, subjects demonstrated at least some degree of reconstruction of their memories: that is, they identified some key or particularly salient bits of information and filled in the rest using background knowledge, inferences, and their own personal tastes and interests. As such, even with semantic material that it should be able to neatly explain, S&R seems to be unable to explain the kind of memory process happening in these cases.

4.6 – Conclusion:

Taken together, everything we've seen in this chapter should suggest that the S&R model is not the appropriate one to be using either for the purposes of understanding procedural memory or for understanding memory at large. The next chapter will be dedicated to taking the seeds of progress sewn by views like Bartlett's and determining how and whether they help with our task of understanding procedural memory.

Chapter 5: **(Re)Construction and Simulation Models of Memory**

5.0 - Introduction:

As we saw at the end of the last chapter, the strict preservationist S&R account was already facing significant challenges in the first half of the 20th century, before it even reached its information processing form. Since the 1990s, it has fallen out of favor with psychologists and has been replaced by a reconstructive account, the seeds of which were planted by Bartlett. Today, the reconstructive account and its descendants (most notably, the simulationist accounts) are more empirically supported and version of them seem to be part of a general consensus among many theorists. As we shall see, this consensus has even started migrating into the philosophical sphere within the work of at least some empirically minded memory researchers. Though both the reconstructionist and the simulationist accounts still utilize the language of encoding, storage, cues/traces and retrieval, they significantly modify their usage, and their understandings of these terms, in ways that avoid the some of the biggest pitfalls of the strict S&R account.

However, as we shall see, there are still some serious pitfalls in the application of these views to our goal of understanding procedural memory. Most notably, both reconstructive and simulationist accounts are primarily focused on episodic memory. The transposition of their mechanics to the procedural case creates several problems that are, first, still similar to the strict S&R account and, second, unable to tell us much about what is happening on the procedural level of skill. It is important to note that several key researchers in this field are explicitly aware that their work on episodic memory would not apply to, or be able to say much about, the case of procedural memory (Michaelian,

2016c, largely building on Schacter & Tulving, 1994). Still, I believe that it is helpful for our purposes to work through these transpositions and first, see why they do not capture the procedural case and, second, how this information may point us towards what we should be looking for in an account of procedural memory. Furthermore, it seems that, in order to be able to say that the episodic case does not apply to the procedural, the researchers in question reduce procedural memory to something altogether noncognitive and inert in episodic remembering. To challenge this, I examine research highlighting the significance of the body (including its procedural memory) for declarative memory at large and use this to offer a way forward into the next sections of this project.

5.1 – ‘Classic’ Reconstruction Views of Memory:

5.1.1 – Reconstruction in Psychology:

As we previously saw, a key insight in (Bartlett, 1995) work is that very little memory is actually exactly ‘reproduced’ and relies, instead, on some ‘core’ element or ‘stereotype’ of an event being retained around which the rest can be reconstructed at the time of recollection. On the basis of his results, Bartlett ultimately identified a few key patterns which he used to help piece together an account of reconstructive memory. For instance, he noted that the manner in which a ‘core’ element of an image or story was identified as well as how details were reconstructed depended largely on the recollector’s methods of recall (reliance on visual imagery or a ‘guiding idea’, etc.), attitudes (confidence or uncertainty, like or dislike, etc.), and rationalizations (related to personal interests, affinities etc). Let us examine some of his results a bit further to see how these ideas were fleshed out.

To begin with, the results Bartlett derived from the first experiment described in the previous chapter (Faces of Military) helped explain how a 'core' was identified. He noticed that that special features of the images tended to be highlighted in the subject's minds and made those features over-potent or over-dominant. This affected recall in ways that showed likely groupings of features, usually around some central detail, which in turn suggested a constructive process of recollection. That is, some central detail or grouping was maintained, and the rest of the detail was constructed around it on the basis of whatever the subject deemed likely to be relevant to the central detail (regardless of whether that was actually present in the original image). Often, the 'remainder' of the image constructed around the core either introduced external details (importation) or moved details from one part of the series to another (transference). This importation or transference happened throughout the various sessions with subjects but was more common in later stages.

Bartlett also noticed that the attitudes people took up during their responses proved to be quite significant for their manner of reconstruction, both in regard to what they noticed about the images and the manner in which they delivered their reconstructions. More specifically, the confidence of visual recollectors and the uncertainty of verbal recollectors, in both cases, had transformative effects on the recalled memories. Bartlett suggests that, usually, subjects didn't take in the details of scene or event, but rather formed some general attitude about them such as like or dislike, confidence or suspicion, amusement or gravity, etc. which then influenced the subject's recollection. For example, if a subject felt a face was 'rather grave', when asked later,

they would often start with some central, core, detail about the face (say, that it was that of an officer) as well as the grave sentiment about it. From there, they would proceed to fill in the details of a grave face (for instance, ‘a stern, older officer with a steely look and a serious expression’).

Furthermore, whatever attitude a subject took up about the images would reinforce or justify the attitude they took up in reconstruction. For instance, in our case, an attitude of confidence that they’d got the face right often led subjects to transform the recollection along whatever path they started – usually, then, their recollections were rather detailed, but not usually accurate. Similarly, for those expressed uncertainty about their recollections, their uncertainty led them to recall relatively bare images without much detail. But whatever was recalled was usually talked through given some core central feature and reported to the researcher. In this way also, their initial uncertainty led to rather bare images which reinforced their uncertain position.

The case of the second experiment (Reproduction of a Story) also brought to light some key patterns. In all cases, the subjects’ reproduced stories demonstrated some level of rationalization, or transformation of the story that made it more ‘satisfying’ or easily understandable to them. Bartlett explains that rationalization happened in a few ways across the subjects but was most commonly seen in the details of the reproductions. Subjects often rationalized the story by drawing new connections between details that weren’t there originally. This served to make the story more readily understandable to them or helped fill some perceived gap. For example, some subjects rationalized a part of the story they didn’t understand by connecting it with some kind of explanation, and

eventually only remembered the version of the story that included their own explanation or elaboration. Similarly, some subjects changed aspects of the story either to be in line with their own personal interests or to make it fit in more neatly with their current time and social group. For instance, some subjects included or highlighted aspects of the story that were relevant to their own professional work, and some changed names and phrases in the story to make them more contemporaneously English sounding.

This particular trend in reconstruction – in which subjects fill in details about past experience on the basis of their current feelings and circumstances - has come to be known as an egocentric bias in reconstruction. As reconstruction approaches became more commonly taken up in psychology, this egocentric bias and its influence has been studied and reproduced across a wide variety of remembered content, including emotions (Levine, 1997), pains (Eich et al., 1985), attitudes (Goethals & Reckman, 1973), and perceptions (Loftus, 1974). For instance, in the Levine (1997) piece, the researchers began by asking supporters of Ross Perot about their feelings (whether they felt surprised, sad, angry or hopeful) soon after he pulled out of the 1992 presidential race. A few months later, when Perot rejoined the presidential race and seemed to be polling well, these same supporters were asked about their feelings immediately following his withdrawal. Levine found that subjects' responses tended to vary from their initial responses and that these distortions systematically depended on their current feelings about Perot and the race. For instance, "supporters who remained loyal to Perot underestimated how sad and angry they had reported feeling after Perot first withdrew from the race and overestimated initial reports of hope" (ibid, 174). On the other hand,

“those who left Perot’s camp and never returned gave stable reports of past feeling of anger, but underestimated initial reports of both sadness and hope” (ibid). As such, these results help bolster Bartlett’s initial findings about the manner in which current circumstances can be brought to bear on one’s recollections and the reconstructive nature of memory.

Just as Bartlett’s work laid the foundations for further research on egocentric bias, it also helped bring to light what came to be known as two central categories of memory distortions: namely, false recognition on the one hand, and intrusions and confabulations on the other. False recognition denotes instances in which people are presented with a novel fact or event that they have never encountered and yet claim to either recognize it or be familiar with it. Similarly, intrusions and confabulations denote, first, instances in which people bring irrelevant information to bear in current questioning and, second, narratively describe events that never happened. Again, the widespread nature of these mistakes, and the continual ability of researchers to produce and induce these mistakes, reinforced the failings of a strict S&R account as well as helped to develop more modern versions of reconstructive theories of memory. Indeed, it became the aim of many contemporary reconstructive accounts to be able to account for these kinds of distortions within their theoretical framework⁵⁴.

⁵⁴ Though this is the explicit aim of constructive views, they have been criticized for being unable to explain the origins and functional differences between different kinds of memory distortions (S. Robins, 2016b, 2016a). However, convincing responses to these critiques have been made (see (Michaelian, 2016b) which suggest the critiques misunderstand certain aspects of constructive accounts, especially simulationist versions (see below).

A key developer of contemporary reconstruction accounts is Daniel Schacter whose career spans many years and many projects (Schacter, 2019). In what we might call his ‘second stage’ of research, he and his lab focused on identifying the neural underpinning of false recognition. After conducting many studies utilizing PET and fMRI data from subjects both with (Schacter, Verfaellie, et al., 1996) and without (Cabeza et al., 2001; Schacter, Reiman, et al., 1996) brain damage, and across a wide age range (Norman & Schacter, 1997), they identified the following key takeaways: that the medial temporal lobe seems key for encoding and retrieving information, that the prefrontal cortex seems key for ‘retrieval-monitoring processes’, and that damage to these areas, either through age or accidents, significantly increases one’s rates of false recognition.

Pulling together much of this research, (Schacter et al., 1998) eventually put together a view they dubbed a Constructive Memory Framework or CMF. They ultimately explain that the tendencies towards both false recognition and intrusion/confabulation relates to the recollectors’ general primary attentiveness to ‘gists’ or general features of events. That is, as Bartlett originally suggested, people tend to retain ‘big picture’ or general features of whatever they encounter, rather than focusing on their details or distinctive features. This means that, when it comes time to retrieve information, they only have very bare frameworks around which to fill in details.

As part of the CMF, Schacter et al. offer the following two points of explanation or clarification about how these memory distortions come about. First, that false recognition is the result of a failure of pattern separation: that is, that patterns for

individual events or words are not sufficiently separated from others, especially those that are closely related or similar to each other, such that the patterns bleed into each other and only thematic “gists” are maintained. For example, if one is being asked to remember events that were sufficiently similar to each other (say, baking lemon cookies on the one hand and baking lemon cakes on the other), it might be that the two events bleed into each other such that they lose some of their distinctive qualities and primarily maintain a ‘gist’ or overarching structure such as ‘baking lemon-flavored goods’. In trying to recall one or the other event, then, one might call forward this gist rather than specifics of either event.

Second, CMF could draw on the notion of ‘implicit associative responses’, or “the idea that people overtly or covertly generate a nonpresented lure word at the time of study in response to an associate” (ibid, 295). This would amount to a kind of ‘source confusion’, where subjects are unable to recollect whether the retrieved “memory” is one they actually experienced or one they generated spontaneously. For instance, suppose one was studying a list of words about baking and, while doing so, thought of some words related to cooking. When asked to recall the words one had previously studied, if one recalls words about cooking, they might be unable to determine whether those words were actually part of the list they studied or just an association they had made while doing so. As such, there would be some confusion as to where the cooking words had come from – actual experience or mere association.

To counter these tendencies, Schacter et al. highlight the following two potential remedies: first, focusing on pattern separation at the encoding phase – meaning, while

one is performing some action or event (or shortly afterwards), one should focus on noticing details and potential overlaps and pitfalls for one's recollection of events. This way, one can minimize the kind of pattern recognition and source confusion mistakes described above. Second, one should utilize focused and specific descriptions or cues at the time of retrieval to help pinpoint the target information – meaning, one should try to use targeted and specific words, cues, etc. at the time of recollection that are most likely to bring forward the correct memories.

5.1.2 – (Re)Construction in Philosophy:

As the constructive view became more prominent and, indeed, eclipsed the more standard storage and retrieval view within psychology, some philosophers began to take notice and tried to work through its implications for philosophical work. This work covers a wide range of issues including arguing for the significance of reconstructive accounts for work on extended theories of cognition and on collective memories (Sutton, n.d.). However, for our purposes, I will focus on a prominent figure working through these implications for epistemological understandings of memory, namely, Kourken Michaelian, who put forward a “generative” account of memory (Michaelian, 2011). Since the reconstructive account is based in the idea of re-creating a memory around a core framework, Michaelian noted that it would entail a generation of new thoughts and beliefs either before information is retrieved (sometimes broadly called construction) or through the process of reconstruction during retrieval. In the former case, this could be either in the process of consolidating a ‘core’ or ‘trace’ – suggesting that even this is constructed – or through some modulation of a ‘core’ or ‘trace’ before it is reconstructed.

In short, every stage of storage or retrieval could entail some generation of new content. If or when that newly generated information is accepted at the time of retrieval, this would also entail new beliefs had been accepted by the subject. Beliefs about remembered content will then be based on both stored and constructed content.

This being the case, the question of the reliability of memory might be raised, especially given the variety of memory distortions discussed above and their frequency. In response, Michaelian suggests that, under usual circumstances, the generation of new content is in line with its original sources, but the process is not infallible and is open to deviations. Indeed, harkening back to the same research on the adaptability of reconstructive memory suggested by Schacter et al., he suggests that this adaptable nature, and the influence of whatever experiences impact the transformations it makes during construction and reconstruction, are meant to ensure reliability (ibid, 330). Consider this his ‘reliabilist’ understanding of constructive memory.

With this in hand, Michaelian revisits a classic piece of philosophical work on memory, namely (Martin & Deutscher, 1966), which puts forward a causal account of memory. In general, theirs is a view suggesting that there must be a particular kind of causal link between a memory one stores or retrieves and an actual experience. More specifically, they suggest, first, that there must be a direct causal chain between the original impressions and the retrieved content via a memory trace and, second, that the memory trace must be active at the time of retrieval and directly contributing to the process of remembering. In updating their arguments in line with new, reconstructive understandings, Michaelian suggests two main modifications: first, that a new version

should include the aforementioned reliabilist understandings and second, that the relation between the initial impressions, memory traces, and retrieved content should be limited in a such a way that the content of traces does not go too far astray from the content of original impressions, and the content of retrieval does not go too far astray from the content of the trace (ibid, 333). In short, he suggests that memory should be understood as generating significant new content but that there must be important limits on it for the generated content to be understood as remembered.

All of this has implications for epistemological questions about the justification of memory. Within the classical preservationist understandings, memory is a belief dependent process (e.g. one that takes beliefs as inputs and delivers the same belief as an output). As such, memory does not generate, but only preserves, the justification of beliefs. However, on a generative account like Michaelian's, there are several possibilities to think through. There may, indeed, be cases where memory is belief-dependent and merely preserving the justification of beliefs. However, on the generative account, memory can generate new content as well as new beliefs based on new content. As such, it can also generate justification for both belief-dependent and belief-independent memories. In both cases, whatever new content is generated may justify the new beliefs it generates along with it (ibid, 337).

5.2 – Simulation Accounts of Memory:

5.2.1 – Simulation in Psychological Literature:

As reconstruction views became more common place within psychological literature and was more thoroughly investigated, connections to other cognitive capacities

were discovered. Again, Schacter and his lab were key players in this development and worked on what would ultimately be known as a simulationist understanding of constructive memory. This began with a study demonstrating similar patterns of activation in both past and future event tasks across what is thought to be the ‘default mode network’ (which includes the medial temporal and frontal lobes, posterior cingulate and retrosplenial cortices, and lateral parietal and temporal areas) (Addis et al., 2007) (for a review about the ‘default mode network’ see (Buckner et al., 2008)). Ultimately, Schacter and his lab suggested that these regions should be understood as a core underpinning of the process of remembering past experiences and imagining future experiences (Schacter et al., 2007).

On the basis of these results, Schacter and his lab put together a set of ideas they call the *constructive episodic simulation hypothesis*. They theorized that episodic memory is not limited to remembering the past, but also includes a capacity for flexibly recombining past experiences in imagination in order to simulate possible future experiences. In other words, they suggested that, since both remembering the past and imagining the future rely on nearly identical parts of the brain, episodic memory could not be limited to remembering the past but must also include the capacity to imagine the future. Indeed, the manner in which one imagines the future relies deeply on one’s memory of the past – those imaginings are constructed from ‘copy and pasted’ elements of past experience, if you will, that are recombined in new ways to create and ‘see’ future scenarios. As such, one can simulate the future through recombining elements of one’s remembered past.

This overlap between remembering the past and imagining the future could account for the kinds of memory distortions and errors discussed above: if one were to mistake things that one merely imagined for things one ‘truly’ experienced, one might fall prey to false recollections, intrusions or confabulations. However, though this flexibility opens one up to these kinds of memory distortions, Schacter and his lab suggest that this should not be understood as inherently problematic, but as an adaptive mechanism. This overlap means that, rather than having to act out potential plans or ways of acting, one can use one’s memories to do this via a simulation. (Schacter & Addis, 2007b, 2007a)⁵⁵. Though it can lead one astray in some instances, the ability to create these simulations ultimately opens up many more helpful possibilities than unhelpful ones, and means that we can save a great deal of time, effort and energy.

Indeed, much ink has been spilled on empirically substantiating the neural similarities between remembering the past and imagining the future. More recently, this simulation similarity between them has been explored in a number of different dimensions, including the nature of scene construction (Mullally & Maguire, 2014) and the relation between remembering the past, imagining the future and counterfactual thought (De Brigard, 2014a; De Brigard et al., 2015; Goldman, 1992)⁵⁶. This mass of research has led to a number of researchers to suggest that there is not as strong of a

⁵⁵ (Szpunar et al., 2014) explores the ability of what they call ‘future thinking’ and break it down into a preliminary taxonomy of four basic forms: *simulation* (constructing a mental depiction of the future), *prediction* (estimating the likelihood of some future outcome, one’s reaction to it, etc.), *intention* (setting a future goal) and *planning* (identifying and organizing steps to achieve a goal).

⁵⁶ Indeed, (Williamson, 2007) suggests that the heart of philosophical methodology rests on being able to simulate belief in the antecedent of a counterfactual and being able to work out what would happen in that situation in one’s imagination.

distinction as previously thought between memory and imagination, and to instead suggest that these are iterations of the same capacity for what is called *mental time travel* (Suddendorf & Corballis, 2007).

5.2.2 – Simulation in Philosophical Literature:

Within philosophy, the simulation theory and the important connection between episodic memory and episodic future thought identified in psychological literature, has been taken up more frequently in recent years. For instance, (Shanton & Goldman, 2010) took up this literature to supplement and complement their own work on a simulation theory of mind reading. In it, they note several similarities between the two systems and suggest that “the best explanation of these similarities is that memory and mind reading employ the same basic type of mechanism” (534); namely, simulation. Shanton and Goldman suggest that mindreading another happens on the basis of simulating the others circumstance from one’s own point of view and inferring from that simulation what the other’s mental states might be. Similarly, Shanton and Goldman seem to suggest that remembering is a form of self-mind-reading, whereby one simulates a scenario of past behavior and infers from that what and how one experienced something.

However, a key proponent of the simulation theory of memory within philosophical literature of memory is, again, Kourken Michaelian (see Michaelian, 2016a , 2016b, 2016c.). Michaelian’s view is perhaps a more radical or unintuitive version of simulationist views than others (see De Brigard, 2014a, 2014b; Perrin, 2011 for some less radical versions and Debus, 2014 for some worries or objections to simulationism in general). This is because his views are meant to avoid the kind of preservationist

tendencies maintained in reconstruction and generative accounts, and those that may still be lurking in the psychological understandings of simulation. As such, and because Michaelian's are amongst the most developed versions of the view, I think his work is worth examining more carefully than other iterations, especially for our purposes.

In developing his understanding of the simulational nature of memory, Michaelian seemingly rejects his earlier generative account, deeming it insufficient. Indeed, he says that the constructive/generative account still suggests a broadly preservationist understanding where, despite there being an important sense of distortion or modification at every stage, there is still a sense of representations being essentially preserved in order for it to be remembered. Instead, he suggests that we should see episodic memory as a general constructive system "designed to draw on information originating in past experience to simulate possible episodes" (Michaelian, 2016c, 103). When imagining a future event, one is drawing on a range of past experiences to model a target event. But, on his line of thinking, the same must be true of remembering the past: in remembering the past, one is recombining past experiences to simulate a past event. Usually, this process will reliably draw on experiences based in the event in question, but it may not always do so and does not strictly speaking need to do so.

To explain this, Michaelian uses an example of walking through a train station which has a large statue at its front. As one walks by it, one might encode a short-term representation of that statue (which may or may not be consolidated into long-term store), or one might not encode a representation of it at all. Regardless of whether it is ultimately

encoded or consolidated, one could still be said to remember that statue on the simulation account. Michaelian says,

In neither case need the absence of stored information deriving from my experience of the episode prevent me from remembering it. If sufficient information originating in other sources – for example, stored representations of similar episodes, or semantic knowledge of the layout of the relevant area – is available, my episodic construction system may produce a representation of the episode when I later attempt to remember it. When this occurs, the simulation theory implies, I remember the episode, just as much as I do when part of even all of the content of the retrieved representation originates in my experience of the episode (ibid, 112).

In other words, in order to remember the statue, I needn't have direct experience of the statue as long as I have other relevant past experience that I can use to help simulate it in my recollection of the train station. The relevant past experience could be having seen a map of the train station and the location of the statue, or having seen a picture of it somewhere else, or having seen similar statues in other train stations, etc.

This being the case, one need not necessarily have causal memory chains that pick out specific episodes of experience. There may be cases in which there are causal chains from past experiences to simulated memories, and this may even be the most common case. But cases in which a past event is simulated on the basis of other experiences should, according to Michaelian, be understood as having been remembered in the same sense as those with direct causal chains of experience. Indeed, he says that, on the simulationist reading, one could remember an entire episode that one did not experience “as long as the relevant representation is of an event belonging to one's personal past, and as long as it is produced by a properly functioning episodic construction system” (ibid, 118).

Consider the following: in cases where one is simulating only parts of an episode, such as in our train station case, even if one did not actually experience the presence of the statue at the front of the station, as long as one was simulating it being there on the basis of actual experience in the train station and any other relevant information, one could remember the statue one didn't notice. Similarly, in cases where one is simulating complete false episodes, such as the classic cases of being lost in the mall (Loftus & Pickrell, 1995), as long as one was constructing the simulation on the basis of actual past experiences, we should not understand that as a failure of memory but as a misremembering of those actual past instances. Furthermore, suppose the person being asked to remember an incident of being lost in the mall actually did get lost in a mall, but were too young to remember it. In this case, rather than being 'misremembered', their simulated memory would be genuine remembering, even if they never 'truly' remembered it having happened.

Of course, this then raises the worry of how to distinguish between memories and 'mere' imaginings. To this, Michaelian suggests that the normal functioning of the episodic construction system helps make this distinction. In normal cases, the representations generated by a healthy subject aim at accuracy – though this may go awry in many ways, the system, overall, is still aimed at some kind of faithfulness to reality. However, in the case of 'mere' imaginings, the same system is being used but is not aimed at accuracy – in fact, it may be aimed at inaccuracy. As such, Michaelian suggests that the main mechanism underpinning the difference between these uses is intention -

“either the intention of the subject himself or the ‘intention’ of his episodic construction system” (ibid, 112).

5.3 – Procedural Memory on a Reconstruction Readings:

Starting with the more ‘basic’ understanding of reconstruction, let’s see how that might apply in the case of procedural memory. Recall that the central premise of a reconstruction account is that, rather than exactly reproducing a piece of information or a past experience, one retains (or stores) some core aspect of said experience. How that core is identified will depend on the subject’s interests and attitudes surrounding the experience in question. Later, when the subject is recollecting that experience, they will retrieve this core aspect and will reconstruct the details of that event around that core understanding. The manner in which the details are filled out will also depend on the interests, attitudes, and rationalizations of the subject. So, transferring these ideas to the case of procedural memory would suggest that, for any given skill, there is some ‘core’ that is stored and later retrieved, and details around it can be filled in as needed.

To see how this might work, let’s work through an example. Going back to our case of Megan Rapinoe performing her bent shots on goal, we would need to look for what the ‘core’ of the skill would be that would be stored somewhere in her memory. Suppose that ‘core’ is something like the central ‘procedure’ we identified in the last chapter – something like “start at the ball and take x number of steps back at a stride length of y ; run back up to the ball at pace z ; hit the ball with this precise part of the foot, f , at precisely the right point, p , with precisely s amount of swing”. We would hope, given our understanding of skill from Part 1, that this would be a procedure that includes

declarative and non-declarative, semantic/episodic as well as procedural, memories.

However, it is worth noting that this ‘procedure’, as written, includes primarily semantic memories and whatever the contribution of procedural memories to it would need to be identified. Let us see, then, if the reconstructive account can help us say anything about this.

Back to the case: Rapinoe is practicing set bent penalty shots from the penalty spot and would need to recall her ‘core’ procedure – again, assume it is largely declarative. From there, she would fill in the details based both on past experience and on her current circumstances. As far as filling in general details, one suggestion is that she would be filling in the general (semantic) procedure with specifically procedural memories: beyond just ‘hit the ball with this precise part of the foot’, maybe how/what it feels like to get that precise part of her foot to the ball; beyond ‘hitting the ball with precisely *s* amount of swing’, how/what it takes to achieve that swing, etc. Further, as far as present circumstances go, if Rapinoe is in peak form (not fatigued, say), she may not need to reconstruct all that much around her ‘core’. That is, the use of present circumstances to help fill in the details of the core would be minimal in this case. If she is beginning to get tired or fatigued, she may be able to reconstruct some details around the ‘core’ based on current read of the situation: that is, she may incorporate details to the ‘core’ having to do with how to counteract her fatigue, etc. Similar things would be true in a case where Rapinoe is performing a bent shot on goal during live play. She may retrieve the ‘core’ procedure for bent shots and can reconstruct details around it based on procedural knowledge and everything she sees and anticipates around her, be it incoming

opponents, the expected position of the ball being passed to her, the position of the goalkeeper or the like.

There may be several advantages to this reconstructive approach as it seemingly avoids some of the main pitfalls of the preservationist S&R account described previously: to being with, it seems to avoid ‘frame problem’-type issues as well as needing to remember and delineate individual experiences. On the reconstructive account, it doesn’t seem that one would need to rely on (specific) past memories in order to be able to reconstruct the details around the core. Instead, one could rely on their general or current understandings and inclinations to do so which, though clearly influenced by past experience, need not be drawn from particular instances in the moment of reconstruction. That is, one’s general understanding of the present circumstance, built up through general past experience and current factors, would be enough to fill in the details without having to draw on particular episodes or events.

This being the case, the reconstructive account seems to avoid some of the preservationist S&R account’s problems around responsiveness, as well. That is, whatever one reconstructs around the core would leave a lot of room for being responsive to current circumstances and being able to adapt existing procedures to them. A reconstruction that takes current circumstances into account will seemingly necessarily be responsive (though perhaps to greater or lesser degrees of success). Similarly, it is easier to understand how performance might be different every time it is done on the reconstructive account: since each reconstruction is tailored to present circumstances and each ‘present’ circumstance would be (at least slightly) different, each reconstruction

would also be (slightly) different. In this way, reconstruction seems much better able to account for our RRRR criteria than the preservationist S&R model was. So far, the transposition from episodic to procedural looks promising.

However, since accounting for distortion is central to episodic reconstructive accounts, we might expect to see some manner of distortions in the procedural version, as well. Consider that, in the episodic case, one might reconstruct some of the details of an episode incorrectly because one was focused on a gist rather than details, or had some sort of source confusion. For instance, in an episodic case, suppose I remember the dining room in my friend's house when I went to visit them a few weeks ago. Suppose my 'core' memory of this room was that it had bright yellow walls – I found them very striking and noted to myself how uncommon it is to see bright yellow walls. Later, when another friend asks me about it and I reconstruct the memory of the dining room, I recall the 'core' yellow walls and can fill in all the details about that room. For instance, I might recall that the dining rooms had windows, a dining table and chairs, maybe I recall a shelf with some decorations. As I reconstruct all of this, since I took in the 'gist' of a yellow dining room and did not attend to these latter aspects, I may very well get the details wrong (maybe I misremember the color of the table, the number of chairs, the width of the windows, the decorations on the shelf, etc.). In this case, I may just be filling in the details of the room based on my general understanding of tables, chairs, shelves, etc. When this happens, I may have a generally accurate picture with some minorly incorrect details, or I may have a picture that is generally accurate but starts to include features that

weren't there, or I may have a generally inaccurate picture with only the 'core' remaining somewhat accurate.

One suggestion for a procedural equivalent could be that procedural memories could be distorted in ways that help explain deficiencies in performance. For instance, if one reconstructs the procedural memories around the core and gets some of that application incorrect, one might still have a generally successful reconstruction but fail to perform exactly as one intends. Suppose Rapinoe is, again, reconstructing her bent penalty shot around the 'core' procedure. We suggested before that the reconstructed 'detail' would be procedural understandings of how to do so, or understandings of what it takes to do so, and acting in such a way as to stay in line with, or replicate, those memories. As she adds in the 'detail' about how to get her leg to perform the procedure, she may reconstruct some of those details incorrectly or with insufficient accuracy. As such, she may perform a generally good bent shot on goal but the inaccuracies in procedural 'details' may prevent her from getting the shot exactly right or doing so just as she means to; or, she may manage to get the bent shot off but off-target; or, she may simply fail altogether.

However, it seems strange to think that the procedural 'details' being filled in could be understood as 'inaccurate' in the same sense as episodic 'details'. Indeed, it doesn't seem that non-declarative memories aim at accuracy at all. I can misconstrue the details of a room and, in that sense, be inaccurate, but can I misconstrue the procedural knowledge of 'how to do something' or 'what it takes to get it done'? Of course, one's performances or attempts at a skill can be accurate or inaccurate, one can perform

flawlessly or stumblingly. But this does not seem to say much about the memories that underly the performance or the skill itself. Given that our understanding of skill suggests that it uses both declarative and non-declarative memories, we can see how declarative aspects may be misconstrued or inaccurate, but it's not clear how that could be said of the non-declarative aspects. Even if we suggest that 'how to do something' or 'what it takes' includes some declarative aspects which can be inaccurate, again, we haven't said much about how the non-declarative aspects can be inaccurate.

This problem is related to a deeper one: namely that it is unclear why we should think there is such a thing as a 'core' or a 'detail' of a procedural memory, and that doing so smacks strongly of preservationist S&R undertones. As such, it is not clear that we can genuinely talk about 'reconstructing' procedural memories in the same way as we do for episodic memory. To see this, consider the following: in the original case of reconstructive episodic memory, events or episodes were being reconstructed largely with declarative information, be it 'inferences' from past experience or present circumstances. In the case of remembering my friend's dining room with bright-yellow walls, it seems a very clear cut case of 'reconstruction' in a relatively literal sense of the metaphor. That is, I am piecing my memory of the room back together based on past and present circumstances and based on my general understanding of dining rooms. Indeed, rooms, buildings, or scenes seem to be the natural subject of a 'building' metaphor like 'reconstruction'. It is clear what the 'core' feature of the memory is, and how those 'pieces' or 'details' might be filled in around it.

However, in the case of reconstructive procedural memory, it's not clear why we should think there is such a thing as a 'core' that has 'details' and 'pieces' that are separable from it. For instance, in the case of Rapinoe's bent shot on goal, we would have to think that the procedural memories brought to bear in the moment of performance are 'pieces' of the overarching skill of bent kicks used to reconstruct it. For this to be the case, we would have to, in some sense, separate out the 'pieces' from the core and from each other. That is, in the way that the table and chairs are 'pieces' of one's reconstruction of the yellow dining room built around a central core, one's procedural memories would have to be 'pieces' of one's reconstruction of the skill. But this opens the door to similar issues faced by the preservationist S&R account in the previous chapter – namely, that these procedural 'details' and 'pieces' could be discrete and separable from the core and each other and so potentially cordoned off from the previously mentioned benefits of the holistic 'gist'.

Consider the following: the notion that these procedural details are separable from the 'core' suggests that, in some sense, one could separate out the 'details' of a procedural movement, like a kick or swing, from the 'core'. That is, around some declarative 'core' there would be procedural details which are all, in principle, separable from the 'core' and each other, like the 'details' of the yellow dining room are from each other. Though they may be tied together through a general understanding or past experience, the pieces are still, in principle, separable. For this to play out in the procedural case, it would be like thinking of the procedural memories of a kick or swing are like the frames of a film of the movement – the declarative 'core' would be filled in

with procedural ‘frames’, each feeding into each other and part of a broader understanding, but still separable. If that is the case, then it is possible that this generalized, yet decomposable, understanding reintroduces the possibility of losing track of these ‘pieces’ and being unable to act because a ‘piece’ or ‘frame’ is missing. For instance, failing to perform a kick because a separable ‘piece’ or ‘frame’ of the action was not included in the detail of the reconstruction.

But this would entail a very odd possibility – that one could lose a ‘frame’ or ‘detail’ of an action. It’s unclear, however, how this could play out in reality – that one could, mid-swing of kick, be missing a ‘frame’ of the action and so fail to perform the kick. It seems much more plausible to think that procedural movements, rather than being ‘cores’ filled in with pieces and details, are continuous wholes that include the entire breadth of the action. That is, the best way to avoid these kinds of problems or considerations is not starting from a skeletal version of skill and filling it in, but rather working with a ‘complete’ version of the skill which may be artificially described in skeletal form.

However, understanding where the advantages of a reconstructive approach come from might point us towards a way forward and give us clues about what to look for in an account of procedural memory. Indeed, buried within the reconstructive account may be the seeds of a solution to these problems. The reconstructive account utilized a notion of ‘current’ or ‘general’ understandings to help explain how a core can be reconstructed. Recall, a core can be filled in with other past memories about what those episodes are generally like or what one finds appealing, etc. But what does that amount to? How is

that developed? I suggest that what is being implicitly suggested here is something like a ‘rolling mean’⁵⁷ of experience which explains the benefits of a reconstructive approach.

To build up a general understanding of things based on past experience, one must, in some sense, be accumulating past experience in order to do so. However, as we saw in the case of stricter preservationist S&R views, we shouldn’t want this to amount to preserving every instance, episode, or experience one has – this leads to problems already discussed. Instead, what the reconstructive approach may implicitly suggest is something like the following: that these experiences are somehow distilled such repeated or frequent features are accumulated and kept together. Think of this as a ‘rolling mean’ or ‘rolling average’: a ‘mean’ in the sense that it is a distilled and aggregated version of one’s experiences and ‘rolling’ in the sense that it keeps integrated one’s new experiences into the mean and updating it. As such, one can have a general understanding of what dining rooms are like, for instance, based on all of one’s past experience with them. This can then be used to fill in details of a scene around a ‘core’ – one is bringing all of one’s past experience to bear in a general way that doesn’t require remembering particular incidents or episodes.

This is, I think, the core advantage of the reconstructionist approach. As we saw, though, once this ‘rolling mean’ is fed back into the framework of a details being reconstructed around a ‘core’, we start to introduce the problems for our procedural case.

⁵⁷ I attribute this terminology to Andy Clark, who I believe used it in graduate seminars I took with him between 2013-2014. Though I cannot find any published work in which he uses this exact phrasing, in (A. Clark, 2016) he uses the phrase ‘rolling present’ and I believe means it in a manner similar to the way I am using ‘rolling mean’. Regardless, I prefer the idea of a ‘mean’ being maintained in the phrasing, as I think it highlights the aggregation of experience more clearly than the term ‘present’ does.

However, when this notion of a rolling mean, extracted from the reconstructionist account, is applied to procedural memories, it may also help explain how a procedural memory could be developed and applied as a more wholistic memory. To see this, consider how one learns or acquires a skill over time. As we previously discussed in Chapter 1, this is largely thought to be a process gradual accumulation and refinement over time. Supposing one has a ‘rolling mean’ of a skill, every instance of practice and training would feed into that mean. One could extract from those experiences ‘core’ aspects of one’s performance and feed them into the mean. In general, as one’s performance changes, those changes would be integrated into the mean and the mean would gradually shift. Hopefully, if one improves, the shift in the mean would be an indication of that improvement and lead to gradually better performance in most cases. So, to remember the skill would not require remembering either all past episodes of practice and training or reconstructing details of performance around a ‘core’ – one would just have to bring the mean to bear on performance.

5.4 – Procedural Memory on a Simulation Reading:

Recall that, on the simulationist account, due to shared neural underpinnings between past and future imaginings, it is suggested that these are actually two elements of the same capacity for ‘mental time travel’. In both cases, when one imagines the future or remembers the past, one is thought to be flexibly recombining past experiences so as to simulate the target event. For cases of future thought, this idea seems relatively straightforward – when one is imagining a future scene, one must draw from past experience to piece together a version of what might happen or what it might be like, and

then one could work through plans or behaviors through that constructed future scene based on one's past experience. In these cases, it seems clear that there could not be a causal source for the experience of the future as it is not something that has happened yet. Similarly, in the case of imagining the past, the suggestion is that one would be recombining past experiences around something from one's personal past and simulating how that episode would have/did take place. However, as Michaelian suggests, these simulations may not require any causal chain leading back to 'real' experiences, as preservationist views might suggest. One could truly be said to remember a past event that one did not directly experience, as long as one was pulling from or aiming towards something from one's personal past.

In considering whether this account could help us in explaining procedural memory, let's come back to our trusted example of Rapinoe and her bent shots on goal. Here, too, the future-oriented simulation seems straightforward. Suppose that she's starting to make her attempt at a set penalty shot. Perhaps she, first, puts together a simulation of what she plans to do – maybe something in line with the 'central procedure' we've described before. To do this, she pulls from her past experiences and can recombine them in ways that help her plan the particular shot she wants to make. Similarly, Rapinoe may remember past instances of play or training and simulate them to help guide her future-oriented simulation or decide what to do in live-play. For instance, she may try to recall an instance of having done this particular shot on goal successfully in the past and so simulate a version of it to run-through and see what made it successful.

To do this, she needn't necessarily draw from one particular source and could instead draw from her general experience of having made these kinds of shots in the past.

This kind of account could be advantageous for similar reasons to the 'basic' reconstructionist approach: that is, this kind of future simulation also seems to avoid issues of finding and remembering particular instances of past experiences in order to be able to perform. Rapinoe wouldn't need to remember every instance of play or practice in order to be able to remember, or simulate, past experiences of potential future ones. Instead, she will likely be drawing from general experience or understanding, similar to the more basic case of reconstruction. As such, it is also easier to see how she could perform responsively to a particular environment (since each future simulation could be taking into account all the relevant current circumstances) and also respond differently to each case.

Another advantage of the simulationist account to the procedural case is that it introduces an element of explicit future orientation that hasn't been discussed in the case of skill-based memory. The procedural case, I think, highlights the significance of 'mental time travel' to being able to perform successfully: in real-time performance, it seems very likely that planning and simulations are importantly at play to being able to, for instance, weave through defenders and get to the point where one's teammate will pass back towards oneself. It doesn't seem unreasonable to suggest that, in order to do that, one would have had to work out how this play might work and how to get there, and that using past experience in simulations would be a promising explanation of how to achieve this.

However, so far, it seems that all the simulations would be based on declarative memories – either semantic information about the particular cases or episodic aspects of previous experience. The extent to which this would include any procedural memory is, as yet, unclear. Furthermore, it seems that, once one tries to focus on the procedural case, one encounters some very unintuitive results that run afoul of some of our established criteria of skill.

Consider a case of trying to remember, or simulate, an aspect of a skill acquired in the past. If Michaelian is right, then one wouldn't necessarily need to have a causal chain leading back to a source of that skill. Whatever aspect of a past skill one simulates, if not based in direct past experience, could either fill in the details of something one had only partially experienced or could simulate something that one had never done at all. However, it seems odd to think about merely simulating a procedural memory actually equating to having that skill. For instance, if I simulate a memory of having played tennis as a child (which I only did very rarely) and focus only on the procedural aspects of that skill (of which I had very little), it wouldn't seem that this 'memory' would equate to any skill of mine. I might genuinely simulate an episode from my personal past, and that simulation may include some genuine semantic knowledge about how to serve, for instance, but, even if this was brought to bear on a present circumstance of play, it doesn't seem that any procedural skill comes along with it. It seems that, in order to have the genuine procedural memory of performing a skill, one would have to have actual experience in that skill built in practice and training over time, as Criterion 1 suggests.

The case of future procedural simulations faces similar problems but could be potentially beneficial for an account of memory of skill. Consider a novice athlete who is simulating the action they are about to perform. Even if our novice puts together an accurate future-oriented simulation (based on the rules they've been learning, etc.), this may have very little to do with how they actually perform. Indeed, it seems common that a novice will deviate from their intended 'plan' because the procedural side of their skill has not yet been developed. However, in the case of a more experienced or expert athlete, future-oriented simulations may prove extremely useful. Indeed, it seems a common tactic among professional athletes to simulate whatever they are about to perform to help guide their actions and set them up for successful execution. In this sense, being able to put together and utilize successful simulations may also be an indication of the development of one's skill. That is, as one becomes more skilled, it may be that one's simulations become more sophisticated, take more into account, or can include more possibilities. This may be in line with what (H. Dreyfus & Dreyfus, 1986) suggest in terms of experts being able to literally see different possibilities and means of completion than a novice, and being able to integrate their past experience more fully onto their environment. Nevertheless, even with this potential link between skilled performance and simulation, if these simulations are merely declarative, we have not yet said anything about the nature of procedural memory and its role in skill.

Indeed, in (Michaelian, 2016c), Michaelian suggests that whatever is said about declarative memory cannot be said about non-declarative memory (including procedural memory) because of several key differences between them. Among these differences, he

explains that, unlike the case of declarative memory, an appeal to stored “rule-based” or “data-based” information, or to information at all, is both redundant and unparsimonious in the case of non-declarative memory. Because he thinks that “in general, a system need not represent the rules governing a learned behavior in order for its behavior to be governed by them” and so does not think that these rules are ever “processed” for the relevant behavior, Michaelian suggests that positing stored information in non-declarative memory would be causally, and so also explanatorily, inert (ibid, 27). Instead, he says that “an appeal to changes at the neural level is sufficient to explain the relevant changes in the organism’s behavior” and he suggests that there is no reason at all to discuss information processing in the context of non-declarative memory (ibid). Though he agrees that non-declarative memory does involve changes to the brain on the basis of experience, he suggests that these changes have nothing to do with making information available to the subject in the future. So, in this sense, Michaelian does not think that nondeclarative memory should be understood as a kind of cognition.

Even though I agree that models meant to account for declarative memory will not say much about non-declarative and that procedural memory likely does not involve ‘rule-based’ declarative information, Michaelian’s view seems to suggest that procedural memory is reduceable to mere motor and/or perceptual acuity – things explainable by the development and changes to neural connections in the brain – and that whatever memory is involved here does not convey any information to the subject worthy of being called ‘cognitive’. However, the fact that procedural memory is not accounted for within declarative models, and that whatever it conveys is not rule- or data-based, does not

suggest that whatever it conveys is ‘sub-cognitive’ or importantly separate from our ‘information-processing’ capacities. On the contrary, I believe that much of what we have explored in our application of reconstructive and simulationist views suggests that their advantages are based in there being important information on the procedural side of the equation not contained in the declarative information and leaves out the role that procedural memory might have in ‘true’ cognitive functioning.

Some of this line of thought is captured by other simulationists resisting some of Michaelian’s conclusions. Since Michaelian offers the account he does at least in part to avoid preservationist problems based in causality, (Perrin, 2018) suggests a way simulationsists could maintain causality in memory without the that the problems that worry Michaelian. The main reason Perrin is concerned with this is because, if we let go of a causal condition (as for instance, Michaelian does), we also lose our ability to explain where the phenomenological experience of a memory being part of our past comes from. And so, rather than giving up on causality in (episodic) memory altogether, he aims to offer a new account of causality based in its procedural aspects.

In short, what Perrin suggests is that the causality we should be interested in is the way in which past experience of construction can cause increased fluency of construction in the future. To explain, he uses the example of constructing two identical jig-saw puzzles: if one was to construct both of them, even though they have identical pieces in identical positions, these two constructions would not be causally identical. That is, the events of construction for each of them would be unique and distinct from the other. However, if one’s performance was faster or more efficient in the construction of the

second puzzle, there would be a causal relation between the first instance and the second instance of construction, namely a procedural one. That is, having constructed an identical puzzle once before, though causally distinct from the second, can still serve as cause of increased performance in the second instance. In this way, repeated or similar construction between the time of encoding the event and the time of recollection can be understood as causally related without needing to posit a causal chain from the ‘core’ memory to the reconstructed memory.

Perrin means this argument to serve a different function, but it also helps illustrate the point that, arguably, even the declarative remembering that Michaelian is interested in could be underpinned by a procedural element. In other words, that even simulation, itself, could be, at root, a function of procedural capacities that help make the simulation effective, efficient, and maybe even possible. A similar argument is advanced by (Hutto, 2017), who argues that autobiographical episodic memory is underwritten by skill: specifically that “the exercise of narrative skills may well be necessary in enabling us to remember our personal pasts in contentful ways” (201). If these lines of thinking are correct, it would suggest, not only that we cannot understand episodic memory without procedural but, more specifically, that the procedural elements make possible what we consider to be ‘information-based’ declarative memory.

5.5 – Conclusion: Potential solutions and a way forward:

So far, we have seen how the reconstructionist and simulationist views have not been able to tell us much about the nature of procedural memory, perhaps largely due to the fact that both are aimed at explaining episodic memory and that their metaphorical

groundings are still misleading. However, both views seemed to have promising aspects that made them more effective than strict preservationist S&R and may have given us some clues as to what to look for in an account of procedural memory. Firstly, the reconstructionist approach gave us a notion of a ‘rolling mean’ of experience that avoids problems of perfect storage and can be used to describe the acquisition and development of skill. The simulationist account, on the other hand, allowed us to see how future-oriented memory may be very important for conversations around skill in general, and around simulation as a marker of the development of skillful performance.

It seems as if what we may be missing is not merely an understanding of declarative and non-declarative memory being involved in skillful performance, but a way of showing their connection and interrelation. To do this, I suggest turning our attention to research around how the body aids in memory processes. As a first stop, research by (Godden & Baddeley, 1975) demonstrated that divers learning material underwater had better memory when tested underwater, while material learned on dry land was recalled better on dry land. This, they suggested, demonstrated a relation between the state of one’s body and one’s ability to remember information.

Similar results have been shown more recently. (Dijkstra et al., 2007) performed a study focused on episodic memory which suggested a link between being in “congruent bodily circumstances” to a target event and being able to remember it more efficiently. For instance, subjects were faster to remember episodes of going to the dentist (and details of those visits) when they were lying down compared to when they were standing up with their hands on their hips. Similarly, they were faster to remember episodes of

playing a sport (which were chosen based on interviews with the subjects before experimentation) when they were imitating the gestures of that sport than when they were seated at a chair with their elbows on the table and their hands on their head. Importantly, all of the ‘congruent circumstances’ Dijkstra et al. discuss often have to do with the position of the body and its (dis)similarity to that of target events. If the position of one’s body can impact one’s declarative recall, this might open the door to the following suggestion in procedural cases: that procedural memory, in so far as it relates to the body and one’s grip on it, may facilitate declarative recall for/about skilled performance.

Similarly, there is some research on gesturing and its influence on memory recall in speech production which seems to make similar suggestions. (Morsella & Krauss, 2002, 2004) put forward a Gestural Feedback Model in which muscle activation related to given semantic information would help retrieve said information. In their 2004 study, they found that subjects gestured more when recalling information from memory than when the target information is physically present before them. The same study suggests that gesturing happens more in cases where spatial working memory is being taxed (for instance, describing something from memory that is either difficult to remember or to verbally encode, such as squiggles or unusual geometric patterns, etc.). In light of this, they suggest that gesturing is especially helpful for recalling spatial information. However, they also found that gesturing happened often in cases where the object in question was physically present and interpreted this as suggesting that gesturing is not only helpful for spatial recall but also for lexical recall. Furthermore, regardless of whether the object was present or absent, they found that restrictions on gesturing

resulted in significantly slower speech rates, suggesting that gesturing taps into sensorimotor features of declarative information and thusly aids recollection and general speech production. Again, what this research might indicate for the procedural case is that, in so far as bodily engagement is related to procedural memory, it may facilitate declarative recall.

Similarly, research by (S. D. Kelly & Barr, 1999) suggests that gestures are so integral to our recollection of speech and its meaning that it is often woven into the subjects' recollection of the speech, itself, which did not include the information conveyed by gesture. For instance, subjects were presented with a video of an actress saying that her brother went to the gym while simultaneously gesturing as if she was shooting a basketball. When subjects were later asked to report verbatim what she said, they were much more likely to report her having said that her brother went to play basketball, thereby integrating the gestural information into the meaning of what was actually said (i.e. that he had gone to the gym). This would suggest that, rather than merely facilitating recall, gesturing is integral to our understanding of what is said in speech and how we remember it.

Though, in general, it is well noted that gesturing aids recall, especially as it relates to speech production, there are a few key theories at play as to why this is, of which I will discuss only a few. These readings tend to discount the role of the body and its significance, focusing more on the declarative knowledge contained within the gestures. To begin with, there is the lexical retrieval hypothesis which suggests that gestures that are semantically related to their targets facilitate lexical retrieval. There is

also the cognitive load account, which suggests that gesturing is most useful when lexical retrieval is difficult, such as in tip-of-the-tongue (TOT) cases (see Pyers et al., 2021). Lastly, there is a communicative account, whereby gesturing is most common and most helpful in instances where one is talking to another person as opposed to when one is alone and, as such, is primarily a communicative aid (see Holler et al., 2013). For our purposes, I do not think it is a matter of great significance which of these views is right. All of them seem to focus on the significance of gesturing to declarative information but, I suggest, the upshot here is different. The important thing to note for our purposes is the role of the body and its movement to our cognition and the potential benefit of focusing our attention on the role of the body in memory in general, especially procedural memory.

This is precisely what the next part of this project aims to do. The next chapter focuses on embodied accounts of memory, examining several key views, to determine whether and to what extent they reveal anything about procedural memory. From there, we will be in a good position to give a final account of the nature of procedural memory.

Part 3:

Body, Time and Memory

The conclusions of the last section of this project have brought us to a point where, in order to give an account of procedural memory, we would need to begin thinking about the relationship between the body and memory rather than treating memory as purely mental, neuronal, or information-based phenomenon. As such, in Chapter 6, we will begin by exploring existing embodied views of memory: first views of memory in general and then more specific views about the nature of procedural memory. As we will see, these views aim at accounts of cognition based in an interrelation between the subject and their environment with the body being the central anchoring between the two domains. For these views, one must always account for the context, environment, and history of the subject in order to describe their cognitive processes. This being the case, processes like memory are rarely understood to be explicable purely in terms of the inner workings of a brain.

However, we will see that having an embodied perspective does not immunize these views from the kinds of mistakes we saw in previous chapters. Indeed, some embodied views explicitly take up general understandings of memory (either preservationist S&R or simulationist readings) that we have already established are not helpful for our purposes. Indeed, one problem with these views seems to be that they take procedural memory for granted and so are unable to explain it. However, other strands of embodied views, especially enactivist strands, tend to focus on an altogether different understanding of memory which they believe is well placed to explain not only memory

in general but procedural memory in particular. There is part of a tradition of seeing memory less as a capacity for preserving information within larger acts, and more as a complete action or doing in its own right. Even so, we will see that, despite believing that their views are natural fits for descriptions of procedural memory, these views actually do relatively little to account for the body, itself, in procedural memory. Indeed, their descriptions tend towards understandings of procedural memory as ‘unsophisticated’ and seem, also, to take it for granted rather than explaining its inner workings.

Chapter 7, then, begins by taking stock of the strengths of the accounts in Chapter 6, as well as their weaknesses. Based on this, I suggest that a few key aspects to an account of procedural memory are missing: namely, a more robust conception of kinesthesia in memory and a more stronger conception of a subject's relation to time and its role in memory processes. To do this, I begin by introducing Maxine Sheets-Johnstone's view of kinesthetic memory and its implications for a view of procedural memory. Further, I suggest that incorporation of figures like Merleau-Ponty who, not only focused on the role of the body in cognition and performance, but also on a conception of time (building on that of Husserl) in which temporal positions and possibilities are a feature of a lived body. Though many contemporary embodied views aim to incorporate similar broad understandings of time into their positions, these two dimensions, I argue, are often lacking in their descriptions of its relation to memory.

With this in place, the way towards an account of procedural memory based around the body in time becomes available. To help facilitate this discussion, I turn my attention back to the most basic kinds of skills, such as walking, and how procedural

memory, properly understood, would happen at this level. I ultimately argue that procedural memory is then neither a question of knowledge conserved or reconstructed, but rather embedded in one's most fundamental experience of the world and one's own body. As such, procedural memory can be understood as an extremely fundamental form of memory and experience that may underlie all others (as well as several other important cognitive functions).

Chapter 6

Embodied Views of Memory

6.0 - Introduction:

The implications of the last chapters suggested that, rather than aiming to isolate procedural memory from declarative memory, it would be useful for our purposes to get a better sense of how the body acts as a bridge between them and how procedural memories can begin to come into play in the performance of skill. Though there are numerous versions and variations of embodied accounts of cognition, from various traditions. For instance, many contemporary views have roots in the phenomenological tradition of Bergson and Merleau-Ponty (Bergson, 2014; Merleau-Ponty, 2012). Some of their descendants arguably include Dynamical Systems Theories (DST) (Esther Thelen, 2005; Esther Thelen & Smith, 1996), sensorimotor views (Noe, 2004) and enactivist views (Hutto & Myin, 2013, 2017). Despite their differences, all of these views operate from a certain, fundamental core: rather than understanding the brain as central to cognition, and mental processes as primarily computational progressions of information processing, they see the body and its engagements as the core of cognition. As (Esther Thelen, 2000) explains in an oft-cited passage,

To say that cognition is embodied means that it arises from bodily interactions with the world and is continually meshed with them. From this point of view, therefore, cognition depends on the kinds of experiences that come from having a body with particular perceptual and motor capabilities that are inseparably linked and that together form the matrix within which reasoning, memory, emotion, language, and all other aspects of mental life are embedded (ibid, 5). This general approach has yielded further tenets of embodied views.

Beyond being embodied (or primarily based in wholistically understood bodies), many suggest we should see bodies as embedded (tightly woven into a given context, environment or history), extended (pushing the boundaries of cognition beyond the body itself), and enactive (where all cognition is dependent on the activity of the body in an environment). Some have dubbed this quadfecta the 4E approach (though each element is endorsed to varying to degrees by various views).

When it comes to understandings of memory within this larger cognitive matrix, there are views endorsing some or all aspects of the 4E conception. Within the ‘embedded’ element, a common thread across many embodied accounts is the suggestion that memory is heavily influenced by external or environmental factors. This is thought to be backed by empirical research including some of the views discussed at the end of the last chapter, and extending to studies such as (Ren & Argote, 2011; Sparrow et al., 2011; Wegner et al., 1991) showing the significance of other people for our individual memory systems. Similarly, philosophical embodied views have been some of the main push for understandings of cognition, and memory in particular, as extended beyond the body (as we will soon see in more detail). Ultimately, and depending on the degree to which they endorse the notion of representation in cognition, many embodied views tend towards understandings of memory, in general, not as the preservation of information, but as a complete action or doing in itself. Indeed, philosophers have picked up on and endorsed a more recent historical push within psychological literature for a very similar line of thinking (Ben-Ze’ev, 1992, 1996; Ulric Neisser, 1996; Shanon, 1998; Stern, 1991).

For our purposes, then, embodied views seem to fit the bill for understandings that link the body, and its role in memory, to all manner of performance and behavior, especially skill. However, because the array of embodied views is so wide and varied, this chapter will only deal with those that, as far as I know, have explicit accounts of memory: either of memory in general, or procedural memory, in particular. More specifically, we will be exploring theories of extended mind, predictive processing, and radical enactivist accounts. In general, embodied views tend to focus on other aspects of cognition, like perception, and do not have explicit views about the nature of memory. This may be because they simply are not focusing on memory, or because they believe their views on memory can be implicitly derived on the basis of their existing account and need not be drawn out. However, by exploring these explicit accounts of memory, we can see that there are important reasons to focus on memory for any embodied account and that it can be a source of either great strength or weakness for the view at large.

Before moving on to more specific accounts of memory in question, I think it is worth reviewing a general framework of what an embodied account of memory might want to strive for. In a survey piece on the nature of embodied memory (in general), (Sutton & Williamson, 2014), make a point of noting what they think is a key component of memory in any action. Namely, that declarative memory (especially episodic) and procedural memory must have a strong, mutually reinforcing relationship in both everyday behavior and skilled performance for any normal function to occur. As part of an illustration of this point, they present an example surrounding Diego Maradona playing at the 1986 World Cup Quarter Final against England, as told by one of his

teammates, Jorje Valdano. As Valdano tells it, in the lead up to their second goal against England, he and Maradona ran through the defense together with the implicit understanding being that Maradona would pass to Valdano to score. As they drew closer to the goal, however, Maradona took the shot himself. He later apologized to Valdano explaining that he suddenly remembered having taken a very similar shot against an English player seven years earlier at Wembley. In the moment, he realized why his previous attempt had failed and what he had to do to correct it – so, he did, and successfully took the shot himself. What Sutton and Williamson take from this example is the following:

Obviously Maradona's success here relies on exquisitely-honed bodily and technical skills. But he was also set apart from other players by effective decision-making at unimaginable speed: here his choice draws directly on a specific past experience. Such precise use of episodic memory to dictate or sculpt present action is striking because it demonstrates the openness of our motor processes to fast, effortless top-down influence, and the conversion of personal memory into usable online form to feed bodily skill in real time. Such interanimation of skill or habit memory with personal memory is not always easy to notice: but, we suggest, it is a widespread feature of the practice of everyday life. (ibid, 5)

Thus far, this description seems in line with our existing understandings of the interplay between declarative and non-declarative memory in skillful performance and should not come as a surprise to us.

Turning our attention back to procedural memory, however, Sutton and Williamson highlight a few more key points about the nature of embodied memory of skill as they understand it – again, most of which should be familiar to us, by now. First, they note that these memories are not rote processes resulting in identically reproduced behavior, but rather flexible and responsive adaptations to every circumstance in which

they are utilized. Similarly, they note that procedural memory is usually very difficult to verbalize or articulate, and that this often causes theorists, like Dreyfus, to draw distinctions between more ‘cognitive’ or ‘minded’ forms of memory and skill memory. Indeed, coming back to some of the work in Chapter 1, they highlight the debate between automaticity and cognitive control and suggest that even though one must often put aside many (distracting) thoughts in order to perform, this does not mean one is unconsciously or ‘mindlessly’ acting in skilled performance⁵⁸. They warn that “it is a mistake to treat embodied memory as so entirely intuitive as to be outside the psychological realm, for this is to reinforce dichotomies between acting and thinking, and between body and mind, which need to be thoroughly dismantled to achieve better understanding of these complex phenomena” (ibid, 6).

This seems straightforwardly to be part of a larger trend in some circles of embodied memory research whereby an attempt is made to bridge traditional distinctions, like body and mind, automatic or controlled, etc. in favor of some more holistic reading of memory. Rather than focusing on the individual contributions of each facet of memory to performance, the push seems to be towards focusing on their dynamic interrelation and mutual reinforcement. Indeed, we have also recently noted at least one theory of how this interplay could happen in cases of reconstruction/simulation, whereby procedural memory may in fact underwrite some reconstructive/simulationist capacities in episodic memory (see Perrin, 2018).

⁵⁸ Sutton makes a similar points in another piece on how the movie *Memento* demonstrates the significant interplay between declarative and non-declarative memories and the ‘mindedness’ of skilled behavior (Sutton, 2009)

However, at least part of what Sutton and Williamson seem to be suggesting is that our understanding of procedural memory should not be one that leaves it ‘unminded’ and ‘unintelligent’. This suggests that it is not only the interrelation between declarative and non-declarative memory that makes skill proficient, but that both ‘sides’ of the relation should be contributing some ‘intelligence’ to the equation. However, the worry with this push towards focusing on interrelation is that it has not tended to leave us with anything but an assumption that procedural memory is at play and takes its mechanics for granted. Indeed, we noted a similar problem in the discussion of Perrin’s suggestion. As such, though the more holistic reframing of memory has its advantages and could be useful for our purposes, we must also be wary of it as an explanation of our target phenomenon.

With this general framework in mind, we can begin turning our attention to views that start to flesh out these embodied possibilities more fully and determine what, if anything, they can start to tell us about procedural memory.

6.1. – Clark and Chalmers and the Extended Mind Hypothesis:

In a famous piece, Clark and Chalmers spearheaded the ‘extended’ element of 4E and suggest the possibility that the mind and its processes could include things outside of the body: that is, that the mind could be understood to extend beyond the brain and the body and include other, external elements. This could include tools (like a notebook or one’s phone) as well as one’s environment more generally (Clark & Chalmers, 1998). To get this off the ground, Clark and Chalmers propose a parity principle whereby, if both internal and external factors can occupy the same cognitive functions in behavior, they

should be equally understood as cognitive or as part of one's cognition. Though their paper does not primarily aim to lay out a view of memory, an account of memory is one of the central aspects and examples of their view and betrays a relatively straightforward reliance on preservationist S&R thinking.

Consider their famous example which centers around two people, Inga and Otto. Inga is a healthy individual with normal biological memory functions. In the event that she wants to go to the Museum of Modern Art (MoMA), she may simply remember its location on 53rd St. and start making her way there. Otto, on the other hand, has Alzheimer's disease which makes his biological memory unreliable. However, rather than relying on his biological memory functions, Otto has a notebook that he carries with him at all times. In the notebook are facts, names, addresses, etc. that he may need. When Otto wants to go the MoMA and can't remember the address, he will look it up in his notebook and use that to help him get where he needs to. Clark and Chalmers suggest that there is no in principle reason to suggest that Otto's memory process is importantly different from Inga's: both must reach into a store of information, in this case, about the location of museums they like to visit, and utilize that information in similar ways. The only difference is that one store is biological and internal while the other is artificial and external. Given that they think this case meets their parity principle, Clark and Chalmers suggest that Otto's memory should be understood to extend beyond his biological mind and include his notebook. Furthermore, they propose that, as long as parity is met, it is entirely plausible to suggest that many external factors could and should be understood as

part of our cognition – the fact that they are not contained within our brains need not exclude them from this.

There seem to be clear advantages to incorporating some extended mind considerations to an account of memory. For instance, it certainly does seem to bolster and help explain how one's external environment can contribute to one's mental faculties. For instance, for many people, it seems intuitively true that the notes, alarms and reminders they keep or set up on their phone can be understood as part of their thoughts and memories. Indeed, it seems often to be the case that alarms and reminders are set up because one knows that they're liable to forget without it and takes advantage of the external scaffolding to prevent this. Though, most likely, one does not have any memory deficiency that would prevent one from being able to biologically remember many of these things, it does seem that off-loading some information onto a phone or notebook or what have you facilitates one's ability to go about one's day.

Along the same lines, there may also be an advantage to a view like this for thinking about skillful performance. For instance, rather than needing to maintain all of the information one needs to perform a skill in one's mind, at least some of that could be offloaded to one's environment. For instance, suppose one is playing guitar and makes a note to oneself in the sheet music to play a particular chord with fingering A rather than fingering B. This may mean that, rather than having to actively remind oneself of the appropriate fingering one has predetermined, one can simply see the note as one approaches that section and act accordingly. Overall, this may improve the fluency and

flexibility of one's performance and could plausibly be understood as part of one's cognitive process in doing so.

However, aside from worries about the claim of extension, itself, this view of memory has been roundly criticized. To begin with, the parity principle proposed by C&C has been criticized for taking insufficient account of the important differences between Inga's biological memory and Otto's external memory (see Michaelian, 2012; Sutton, 2010 – though, for Sutton, this only means that the parity cannot be a functional one and extended mind thinking could still work with a parity principle around complementarity). More relevantly for us, despite its intended purpose of incorporating external elements and one's environment into one's cognitive processing, some have criticized the view for remaining *too* internalist. They suggest that the very notion of *extended* cognition suggests that it is primarily an internal function that can be extended outwards. As such, they believe that the extended mind hypothesis doesn't go as far as it could in suggesting the true interrelation between a subject and their environment and the significance of this for said subject's memory system (Myin & Zahidi, 2015). Most importantly for us, though, the view has been criticized for maintaining a rather clear cut and outdated view of memory as stored and retrieved information and, as such, of memory as a largely preservational mechanism for declarative information (Barnier & Sutton, 2008; Hutto & Peeters, 2018; Loader, 2013; Myin & Zahidi, 2015).

Indeed, in building on these criticisms for the case of procedural memory, we can also see that, at least in the kind of example discussed above, the advantage offered by offloading some information and 'extending' one's mind seems most plausible for

semantic information. It is not clear what this view can tell us about the procedural side of skillful performance given its heavy reliance on preservationist S&R framework. As far as internal, biological memory is concerned, it seems that this account would run into similar problems discussed in Chapter 4 about ‘storing’ procedural information. Further, if taken to be an account of procedural memory, this view would also have to explain how procedural memories could be potentially ‘stored’ externally and offloaded in the same way as semantic memories. This, of course, seems even more difficult in the procedural case.

In a response to the original criticisms, (A. Clark, 2008), Clark seems to take into account some criticism about the outmoded understanding of memory and the differences between the case of Inga and Otto, noting that some more nuance in the description could help avoid these worries. Even so, he seems to double down on his and Chalmers’ initial assumptions and suggests that, regardless of developments in memory theory, his and Chalmer’s claim was merely that there be ‘sufficient functional similarity’ between the cases of Inga and Otto and that this shouldn’t change the point about Otto’s notebook functioning as part of his cognition (ibid, 98). In any case, throughout these responses, he seems to maintain the notion of stored and retrieved information as central to memory functions.

As far as our intention of understanding procedural memory goes, and given that C&C’s brief, general account seems to only be relevant to declarative memory, it seems that utilizing this account for our purposes would land us with similar problems to those discussed in Chapter 4 and add additional worries about how procedural memory could

be extended into the world around oneself. Though the extended mind view does incorporate the role of an environment around the subject and may help elucidate how one could off-load declarative memory to aid performance, these critiques (and Clark's insistence on some manner of S&R view) should help turn our attention to other approaches that do not fall prey to the same theoretical problems, especially as they relate to procedural memory.

6.2 – Andy Clark and Predictive Processing

In more recent years, Clark has taken a new tack in understanding cognition which relies on an understanding of the brain as a predictive processing machine (A. Clark, 2016). As Clark explains it, predictive processing accounts “offe[r] an attractive ‘cognitive package deal’ in which perception, understanding, dreaming, memory, and imagination may all emerge as variant expressions of the same underlying mechanistic ploy - the ploy that meets incoming sensory data with matching top-down prediction” (ibid, 107). In short, this approach consists in understanding much of cognition as part of a top-down prediction driven process based on accumulated past experience, or what Clark calls a generative model of the world. On the basis of this model, one can generate predictions about what one will encounter as one goes through any given circumstance. As experience accumulates, one's predictions may change over time or become increasingly fine grained. In general, one can keep moving through an action or environment with fluency as long as one's predictions are consistently realized. However, when incoming sensory data demonstrates that a prediction is not actualized, one will generate an error signal and readjust. In other words, when one is surprised and

encounters something they weren't expecting, the system will take special notice and change its behavior while making new predictions around this change.

A great deal of this predictive processing account depends on memory, as this is effectively the probabilistically driven generative model at the center of the account. To help account for this, Clark presents the predictive interactive multiple-memory system model (PIMMS), put forward by (Henson & Gagnepain, 2010). On this view, the different memory systems at play in action (semantic, episodic and perceptual memory), feed into each other at every stage of storage and retrieval. Though they are all feeding into each other and giving each other information, each memory system has different patterns of information flow that help to differentiate between their respective domains and provide information for the others about what to predict in 'lower levels' of behavior (such as perception and what to expect to see in a scene, say). The PIMMS model suggests that encoding and retrieval are driven by prediction errors, or the results of conflicting predictions and incoming sensory data. In other words, they suggest that much of what gets stored in memory will have to do with instances that stood out to the system (or surprised it), and that retrieval will be similarly influenced by the strength of the original predication error.

Clark, however, notes that this account is 'incomplete and speculative' and suggests that the main take-away from it is the suggestion that,

“the surface appearance of multiple, distinct neural systems subserving different functions (here, different kinds of memory) may be subtly misleading. Rather than a mere motley of different systems, we may confront a web of statistically sensitive mutual influence that combines context with content, and balances specialization against integration. Within that web, moment-by-moment performance depends on the creation and maintenance of task-specific patterns of

effective connectivity (here linking semantic, perceptual and episodic sub-systems)” (A. Clark, 2016, p. 104).

In short, it seems that Clark aims to highlight the manner in which different systems can play into each other while also maintaining their individual functions. This means that, for any given action, many aspects of one’s cognition, including different memory systems, will be in play simultaneously in ways that respect their individual contributions while also being totally contextually and temporally situated. In sum, aspects of perception, recall and imagination will be in play within any act of prediction based on one’s generative model.

In fact, Clark sees this predictive conception of memory as totally compatible with more recent accounts of mental time travel and constructive simulation whereby episodic memory is understood as a larger capacity for recombining facets of past experience into simulations of both past and future experience (described more fully in Chapter 5). He explains that “what seems to be emerging is a view of memory as intimately bound up with constructive processes of neural prediction and (hence) imagination” (ibid, 107) and this is to be expected on a predictive processing account. In so far as it is compatible with his PP view, then, Clark seems to endorse the simulationist reading of episodic memory, namely Michaelian’s version, and its preservationist tendencies.

With all this in hand, let us consider how this PP model might apply to our case of procedural memory. Coming back to our Rapinoe case, it seems that she would have to have a probabilistically driven generative model that could help guide her performance in shooting bent shots on goal. That is, as she lines herself up for a set shot, or as she sees an

opening in a real-time game, she should have some store of accumulated experiences on the basis of which she can make predictions about what she's going to encounter. As long as her predictions are not met with any contradictory incoming stimuli, she should be able to perform her shots quite smoothly. If, however, some sensory input comes along indicating that her predictions were off, she will generate an 'error' message and have to readjust her predictions in line with them.

As Clark describes it, the generative model she would be using to do this is, in fact, not dissimilar to the 'rolling mean' we described in the previous chapter. That is, this model would incorporate many instances of past experience into a larger, amalgamated understanding of those instances, and become probabilistically weighted such that predictions are generated on the most frequently recurring, or most likely, aspects of said instances. So, for instance, multiple instances of practicing set shots on goal could generate a model whereby the most common occurrences in her past experience would shape her predictions and, as long as she's not encountering something odd, she would be performing under 'normal' conditions. In this way, then, this view still maintains the core advantage of the reconstructive approach in not requiring every instance of past experience to be remembered in order to be able to perform.

However, if she does encounter something surprising, this would catch her attention and may change her performance. For instance, if she starts practicing set shots and making predictions based on her accumulated past experience of performing with her usual energy levels, when she takes her kick on the ball, she may be surprised (and encounter a prediction error) if she finds that she's actually quite fatigued and cannot

perform as usual. Encountering this prediction error would mean that, in order to keep practicing, she would need to start generating new predictions around the fatigue-input. To do this, she will likely pull from other aspects of her generative model, based on other past experiences, and make her adjustments. This process of correction could continue as long as she keeps playing.

Bringing in some of the simulationist implications, we can also see how the simulations would go hand in hand with her predictions and will draw from past experience. For instance, if Rapinoe is just lining herself up to take a shot, she may put together a simulation of what she's going to do based on taking elements of her past experience and recombining them into her intended shot. In this way, we can see how the PP and simulationist accounts do truly seem to be compatible and, as such, how the PP account can also offer an element of future-oriented memory. Further, it may be that the PP account can even work to supplement the simulationist reading with the generation of probabilistic predictions guiding the simulations rather than merely relying on a broad conception of 'past experience'.

However, it seems clear that, even in these applied examples, what the PP account is offering is a view of what declarative memory would offer in instances of skilled performance. Though it seems that PP could offer us another means of understanding how skill or expertise functions in performance situations, with its predictions being more refined or fine-grained than those of a novice (similar to what we discussed in the last chapter), it is unclear what the predictions guiding behavior tell us about procedural memory/abilities at all. Indeed, as far as PP is concerned, it seems that memory is an

updating store of declarative memory that either disregards the procedural elements of performance or takes them completely for granted. In other words, the PP account only works because it is taking for granted the mechanisms of procedural memory and its influence on behavior (keeping in mind examples such as those of Perrin and Hutto from the last chapter as well as general considerations about the procedural history of generating predictions and responding to errors). As such, it cannot tell us about how procedural memory works or what its contributions are, and can only tell us about declarative memory.

6.3 – Enactivist Theories of Memory:

As we have seen, even embodied views that maintain preservationist S&R tendencies and focus primarily on declarative memory are unable to help us in our quest for understanding procedural memory. One solution, then, is to abandon any sense of storage and retrieval in memory by introducing ‘contentless’ cognition, including memory functions, that do not rely on the transmission of information for them to work. This is precisely what radical enactivist views have done. Indeed, some key figures endorsing radically enactivist views have recently turned their attention to memory research because of its repeated use as a counterexample to their views. That is, it was unclear to many critics how an account of ‘contentless’ and non-representationalist cognition could account for memory as it was traditionally understood, i.e. as stored and retrieved information. As such, these enactivist views have aimed to offer up views of memory that they think maintain their original lines of thinking. Indeed, they tend to believe that they are particularly well placed to explain procedural memory and tend to

focus more attention on explaining their views on declarative memory. This being the case, I'd like to examine two pieces that lay out radical enactivist views of memory in general and procedural memory in particular. I mean to do this without taking a stand on whether the overall enactivist approach is sound or whether their views of declarative memory are right – I only mean to examine their understandings of procedural and determine how well they serve us given everything we've garnered up till this point.

6.3.1 - Myin and Zahidi on Embodied Memory

In their 2015 paper, (Myin & Zahidi, 2015) aim to put forward an account of remembering that rejects internalism (understandings of cognition as being limited to the internal processes of a brain or body) and what they call traditional views of memory (broadly, views with preservationist S&R tendencies). Instead, they embrace a view where cognition, including memory, is an 'extensive' process moving across bodies and their environments that should not be understood as 'internal' or based purely on recollecting information.

What they suggest, instead, is that for something to be remembered, it must be *entrenched* in relation to the subject's other memories, attitudes, and experience. That is, for something to be remembered, it must be contextually related to events that preceded it, followed it, and to the rest of the subject's personal past. The kind of contextualization they're after could be either explicit or implicit. In the former case, it might be that subject can recall specifics about their personal involvement in something (for instance, remembering meeting someone at an event because they were struck by their perfume); in the latter case, they may recall something on the basis of their habitual or routine

manner of engaging that information (for instance, recalling the route to work based on how one usually goes about the area). In either case, they stress that “entrenchment need not be restricted to relations between the elements of some autonomous ‘memory level’ or the level of explicitly formulated thoughts at which pieces of articulated information are explicitly linked. It concerns the embeddedness in ‘a life’, with its personal history, individual motivations and emotions that are constitutive of memory and recollection” (ibid, 399).

Since they think that memory must be deeply embedded in the subject’s life and cognition, Myin and Zahidi suggest that it must also have anticipatory aspects. That is, as one goes about one’s business, one is often laying the groundwork for being able to remember something later. Again, this may happen explicitly or implicitly: in cases of practice, training, learning, etc. one may be explicitly marking things to remember for the future. But one might also ‘mark’ things to be remembered in more implicit ways based on whatever one finds salient, interesting, unexpected, difficult, etc. without ever explicitly marking it out for oneself as such. In either case, though, they suggest that there is an ‘internal link’ between remembering (i.e. connecting the past to a current moment) and anticipating a potential time at which one would need to remember (i.e. connecting the present to the future). Indeed, they bolster this statement with research around ‘mental time travel’ discussed in the previous chapter.

As such, they suggest that memory, in general, should be understood in much broader terms than previously thought. Rather than being composed of discreet events or bits of information, memory should be understood as a capacity for behavior which

demonstrates previous interactions with the topic/environment, and as being the result of several converging processes and tendencies within a larger 'life'. Whatever one remembers in the future will be the result of whatever one's dispositions, habits, and history made possible at previous a 'present', itself informed by one's past. Because of this, Myin and Zahidi suggest that memory will necessarily be sensitive to context and situation since every 'present' circumstance will play on our existing experience differently. Indeed, they suggest that differences in context and situation are all that is explanatorily necessary for understanding differences in behavior or performance.

Even so, Myin and Zahidi suggest that one can still influence one's circumstance or situation through conscious design – one could gradually change one's context or provide with means of getting through the same general circumstance differently. To explain this, they invoke a famous example from John Haugeland (Haugeland, 1998) in which, as one drives down a road (say, to work every day), one may place signs indicating where to turn, change directions, or the like. As one drives that route repeatedly, one may no longer need to pay attention to the signs as one previously had, and all it would take to get to one's destination would be to start down the established road. In the same way, one can influence many of one's circumstances and so, what one remembers, over time.

All of this means that memory, in Myin and Zahidi's view, must include external factors – factors strictly 'outside' the body or brain. If remembering is always a question of context and situation, then external factors will always be essential in what one remembers. Memory, then, cannot only be about stored and retrieved information

discreetly represented in the mind, but must be *extensive* across the subject, their environment, and time.

However, it is worth noting that their view, as it stands, risks being overly broad and introducing new problems. Specifically, the idea that memory should be understood broadly as the influence of past experience on behavior, without further clarification, opens up the possibility that falling from a tree and breaking one's arm, or becoming an amputee, is 'remembered' by the subject. That is, that their body could be understood to 'remember' the damage they suffered, insofar as it is still 'present' and is influencing their behavior. More would have to be said, on their part, as to why this unintuitive result should or shouldn't be understood as memory on their view.

6.3.2 – Hutto and Peeters on Procedural Memory:

Hutto and Peeters ((Hutto & Peeters, 2018), like Myin and Zahidi, address concerns of internalism and (over-)dependence on notions of storage and retrieval of information in memory. They suggest that most views of memory (including simulationist views and most readings of memory in the extended mind literature) maintain some version of the following two assumptions: first, the Remembered Content Assumption (RCA) which suggests that "remembering primarily involve[s] access to and recovery of content of some kind" (ibid, 99) and, second, the Contentful Contribution Assumption (CCA) which assumes that "the contributions that external resources make to cognitive processes only count as cognitive contributions if they are information or contentful contributions" (ibid, 100). Instead, they champion a radical enactivist approach on which "memories come into being 'on the spot' during acts of recall" and

“remembering [is] a matter of ‘on-the-fly’ construction that can be grounded in structural synaptic changes in the brain as well as other structural changes in the environment without assuming the existence of stored and retrieved contents” (ibid 104).

Indeed, they build on understandings of memory as neural networks being strengthened through activation over time to explain why memory need not be understood in terms of content or information. That is, if memory is just the strengthening of neural connections in the brain and a disposition to reactivate in a given way (based on the manner in which the connections are strengthened), and this process does not require any sense of representation of content or information being ‘stored’, then there seems to be no role for ‘information processing’ or ‘stored content’ in memory. Indeed, contra Michaelian, they seem to suggest that what is explained by ‘bare’ neuronal processes is not non-declarative memory, itself, but the manner in which it can unfold without needing to be representational or ‘stored’.

Ultimately, Hutto and Peeters aim to “explain basic forms of learning and memory entirely in terms of re-enacted know-how”, particularly, embodied procedures that constitute sustained response patterns (which may be triggered by environmental factors). In short, they suggest that memory “entails knowing what to do in familiar circumstances” (ibid, 105). Again, contra Michaelian, they seem to suggest that this non-representational, non-information-based process should be understood as one that is basic to cognition at large. Rather than being reduced to motor- or perceptual-acuity, Hutto and Peeters seem to suggest that non-declarative memory should be understood as an overall

responsiveness of a being that is based in their past experience or, ultimately, “reinitiating familiar pattern[s] of prompted response” (ibid, 108).

6.3.3 – Assessing Enactivist understandings:

Since they describe memory as, fundamentally, a process of ‘reenacted know-how’ and ‘knowing what to do in familiar circumstances’, it seems that these enactivist⁵⁹ views may be privileging procedural memories as the most basic and foundational kind of memory. Indeed, rather than understanding procedural memory as mere motor- or perceptual-acuity, they seem to be building in a level of ‘intelligence’ or responsiveness into the procedural side of things without it relying on any declarative knowledge or information. We might suggest that, in a sense, what these enactivist views seem to be doing is precisely the opposite of what the various stripes of S&R views were doing – rather than trying to understand declarative memory as foundational and reading non-declarative back into it, the enactivists seem to be reading non-declarative memory as foundational and using it to explain the nature of declarative memory. But I think there is reason to believe that what they’re really after is more subtle than that: a reframing of the discussion.

In sum, these enactivist views seem to be pushing for a much broader understanding of memory than we have been using up till now. They are not primarily concerned with explaining the retention of information as central to memory, but rather in

⁵⁹ Since I believe that the views discussed in these two papers are largely synthesized into one view in (Hutto & Myin, 2017), I will take this as a cue to refer to the two views as not only compatible but explicitly aligned. For ease, then, I will take the liberty of referring to the set of authors across the two papers as ‘enactivist’ views, rather than referring to all four or only Hutto and Myin.

understanding memory as the influence of past behavior on experience. With this move, they seem to be aiming to shift the terms of the conversation, and the central topics of investigation, so that they are no longer based in explaining how information is garnered, retained, reused, etc. Instead, they seem to want to shift the focus of the discussion towards understanding what the influence of past experience on behavior is without having to assume anything about the nature of that influence or rely on preconceived notions about it. In doing this, they present an opportunity to refocus our attentions away from the brain and towards a dynamic relationship between a subject, their personal histories, and their environment.

As far as an account of procedural memory goes, this kind of account seems precisely what we were looking for at the beginning of the chapter. Indeed, the complete rejection of information-based understandings helps steer us very clear of any S&R account of memory. Rather than storing and retrieving information, past experiences, parts of past experiences, or anything else, procedural memory, on this view, seems to rely on reactivating established patterns of behavior that have been built up over time. This kind of understanding naturally brings the subject's personal history into the fold and helps lay an explanatory foundation for the patterns being built. And, since this happens completely 'on-the-fly' or 'on-the-spot', the need for either storage or retrieval is similarly made redundant.

However, since the enactivists seem to suggest that these patterns should not be reduced to mere acuities, whatever is reconstructed 'on the spot' in response to a given situation cannot merely be some rote manner of acuity, but should include larger scale

understandings about the circumstance, one's goals and intentions and, therefore, appropriate ways of acting or responding. In this way, too, the enactivist accounts seem to incorporate a forward-looking intentional element, whereby whatever procedural memories are brought to bear in the moment can or will be guided so that they correspond to the moment and towards future points of action.

Putting all of this together in our familiar Rapinoe case, the enactivists might suggest the following: Suppose that Rapinoe is actively practicing a variation on her usual set bent penalty shot. As this is a situation familiar to her based on a great deal of past experience, as Rapinoe lines herself up for the shot she will necessarily bring all this past experience to bear with her on the shot. But this experience is not stored or retrieved, rather brought to bear instantaneously based on being 'triggered' in the right way by all the familiar circumstances around her. Once 'triggered', Rapinoe can follow through a familiar pattern of response (something broadly equivalent to the 'procedure' we discussed in previous chapters) and, given her goal of trying some new variation, add in her modifications 'on the fly'.

Similarly, in the case of Rapinoe taking a bent shot on goal during live-play, the situation might unfold thusly: finding herself in a situation with which she is largely familiar (considering that she has been playing professionally for many years and that she and her teammates have practiced the particular shot they are gearing up towards for some time), Rapinoe will find herself 'triggered' by the familiar circumstance to behave according to previously established response patterns. Again, this is not stored or retrieved information, but something that is brought to bear 'on the spot' through being

appropriately cued and would make allowances for ‘on the fly’ modifications given whatever differences there may be between her current circumstances and those in which she practiced.

On the surface, this may seem a rather appealing presentation, especially given everything we’ve encountered up till now. Even though I am inclined to agree with the overarching approach to this topic and with the broad strokes of their conclusions, I suggest that this account is still largely uninformative for our purposes. Though it gives us many of the things we were looking for (a sense of responsiveness, a future-oriented element, a sense of accumulated history brought to bear in performance), I think the manner in which it does so is largely unsatisfactory.

For instance, the notions of temporality described by the enactivists seem incomplete for accounting for our target phenomenon. Their account does aim to explain how action must have a prospective, future-oriented element that helps determine how and what one remembers. They take this to explain how one might gradually come to adjust one’s behavior patterns or how one might anticipate future needs. They back this up with research on mental time travel and seem to suggest that at least part of what’s happening in these cases is the subject simulating a future need and ‘marking’ something to be remembered on that basis. This, however, does not give us a sense of what is happening in live, on-line performance and how one might need to ‘project’ oneself to the end of a sequence in order to make adjustments to it. Similarly, since they explain one’s relationship with the past as a reactivation of old patterns, they seem to lay out a sense in which the past can be dormant one moment and brought to bear the next. But this

might implicitly suggest that one is acting moment to moment, moving from one triggered activation to the next, rather than explaining how any of these moments hook up in experience or performance. Without that, it's unclear how the subject would connect the various reactivations they experience.

Similarly, consider the nature of the 'triggering' described by the enactivists. Though they do describe the manner in which one can oversee the unfolding of a triggered pattern through one's intentionality, it seems the triggering itself is largely passive and inert. If I've got them right in our Rapinoe case, then it seems that merely finding herself in a familiar circumstance is enough to bring the appropriate learned response bubbling up to the surface. But this seems to leave out any sense in which Rapinoe's goals and intentionality influence not just the unfolding of the pattern, but the triggering of the pattern itself. That is, I presume that Rapinoe does not merely find herself performing the kick and then implementing some oversight, rather that she initiates the appropriate pattern through her own volition. As such, it would seem that the enactivist account of the intentionality of procedural memory is not as fully developed as we would need it to be.

Lastly, though they believe that their view is a natural fit for explaining the nature of procedural memory, the enactivist's view tells us very little about the physical contributions of the body to the nature of procedural memory, especially in its relation to skill. Most of their examples, for instance, have more to do with episodic cases (like recognizing someone from a party) or habitual cases (like driving to work in the usual way). Even though they aim to take the body and its place in an environment into

account, their account tells us very little about the physical contribution of the body towards skillful performance besides finding itself in a given circumstance. Indeed, given their intended aim of suggesting the ‘intelligence’ of procedural understandings and its not being reducible to mere motor- or perceptual-acuity, they still seem only to be able to tell us that procedural memory involve the right kind of habituation, rather than explaining what it is that the body ‘knows’ or is contributing to skillful performance – even a very basic one like walking. If this is right, I believe they are, like Clark, taking procedural memory and its usual description for granted and building an account around it, rather than offering a true explanation of it and its inner-workings. Part of this may be because their view is too vague or underspecified, as is. Because of this, what they primarily offer is a broad picture of what we should want an account of procedural memory to be able to account for, without being able to tell us much more than that.

Given all of this, and despite all the attractive benefits and advances of the enactivist approach for the case of procedural memory, I think there is still important work to be done in order to give a more complete accounting. This is what we will aim to do in the next and final chapter, aiming to account specifically for the faults noted above. However, before moving on to those, I wish to consider one final approach as one that starts us on the path towards the last chapter and offers some remedies to the previous views while also remaining incomplete.

6.4 – Bartlett on Procedural Memory:

Throughout the last few chapters, have kept coming back to this idea of the ‘rolling mean’ of experience and its advantages. Indeed, we initially identified this as the

central strength in applications of reconstructive approaches to procedural cases. We have also kept coming back to notion of temporality in procedural memory. Bartlett, himself, offered an account of procedural memory that takes these concerns into account in a promising way (Bartlett, 1995). Rather than understanding his approach to procedural memory as reconstructive (as some do), I suggest that his view is much closer to, and even precedes, later phenomenological understandings and their readings of movement and bodily action. Indeed, many of his ideas on this transposition do not necessarily rely on the same reconstructive aspects described in the previous chapter: as we shall soon see, his procedural account does not rely on piecing memories together around a core, but bringing them to bear as whole and complete influences on one's behavior.

In beginning his considerations of how to apply his declarative findings to the case of the body and skill, Bartlett introduces a problem of trying to understand how past experiences can influence behavior and, in particular, how past movements can have some regulative function on future ones. Following a course similar to our own investigation, Bartlett indicates that an older version of S&R (based in traces and associations drawn between them) seems a natural and easy starting point to answering this question and works through some contemporary S&R-type understandings of memory of skill. In the more general sense, these understandings rely on reacting to incoming stimuli to pull on relevant stored 'traces' and so incorporating them into current behavior. On the smaller scale, in order to perform an action, these views understand the subject to have to break down a movement into many smaller constitutive parts. These parts are somehow stored and can be 'activated' in the right order, and this gives the

subject the ability to recognize and track the transition between movements as well as larger movements, in general. In this way, one is thought to be able to (exactly) reproduce the same movements over time.

However, Bartlett quickly pivots away from this kind of notion, citing his own experiments that suggest “the past operates as an organized mass rather than as a group of elements each of which retains its specific character” and describes how the movements of skilled behavior seem to operate as a whole, rather than a series of disconnected actions (ibid, 197). To help explain this, Bartlett puts forward a notion of ‘schema’, which refer to,

“(...) an active organization of past reactions, or of past experiences, which must always be supposed to be operating in any well-adapted organic response. That is, whenever there is any order or regularity of behavior, a particular response is possible only because it is related to other similar responses which have been serially organized, yet which operate, not simply as individual members coming one after another, but as a unitary mass” (ibid, 201).

In other words, rather than having a series of stored movements in place that are serially connected to each other but maintained as individual ‘units’, Bartlett suggests that these progressions of movement are kept together as a cohesive whole, a ‘unitary mass’: as ‘schemata’. As such, one’s past operates on current behavior “*en masse*”, with all one’s past experience acting at once: or, more specifically, the schemata are predominant influences on current behavior, rather than absolute determinants.

Though these schemata are based in past experience, they will usually be grouped ‘thematically’, according to Bartlett: that is, they are related back to a given setting or general environment. On a ‘low-level’ the organization of these settings will involve visual, auditory and other sensory experiences; on a ‘higher-level’, larger, common

interests, such as sport, literature, art, science, philosophy, etc. will organize them. In other words, various settings, be they strictly environmental (i.e. the right sort of physical setting) or based around an interest or activity, will be closely associated with particular schemata and will be involved in bringing them to bear on action.

Similarly, various experiences within these ‘groupings’ all contribute to the eventual mass of experience that is a schema. Often these experiences maintain a sense of the chronological order in which they happened but, in many instances, particular occasions or events may be given priority based on salience to the subject. That is, parts of schemata that have strong personal interest, relevance to ideals, appetites, instincts, etc., may have a highlighted place within the schema and so function as a particularly strong ‘trigger’ or ‘trace’ for them (ibid, 210). Indeed, schemata, themselves, might overlap and have different priorities, again depending on the subject’s appetites, instincts, interests, etc⁶⁰.

However, all of these experiences cannot be understood as parts of some passive framework pieced together by the subject, or “as a number of individual events somehow strung together and stored within the organism” (ibid, 201). They must, instead, be “regarded as constituents of living, momentary settings belonging to the organism, or to whatever parts of the organism are concerned in making a response of a given kind” (ibid). That is, the contributing influences of schemata, or the ‘active factors of schematic

⁶⁰ This, if anything, is what Bartlett thinks explains ‘temperament’ and ‘character’: how and what we remember has to do with some ‘special active pattern’ which inevitably takes on idiosyncratic characteristics which, in turn, expresses the character or temperament of the person doing the remembering.

organization', should be understood as 'living', changing, and adapting phenomenon that are responsive to particular contexts and environments, rather than set or static. This means that different schemata might be most relevant even in overarchingly similar contexts: depending on the setting the subject finds themselves within and their interests, tendencies, and appetites in the moment, different responses may be called for even in very similar circumstances.

Furthermore, Bartlett brings in a temporal dimension to the continuity of movement and how the subject can keep track of them in performance. He says that, as we move, our actions are related back to the 'schema' in play – as soon as the relation is made, postural recognition (i.e. understanding of the location of one's limbs, characterization of this new bodily posture, etc.) follows. Past experience, then, doesn't rise into our consciousness directly, but sets a foundation against which to measure all current action and movement and keep track of it. This is how he explains the fact that in every skilled performance, "a large number of movements are made in succession, and every movement is carried out as if the position reached by the moving limbs in the last preceding stage were somehow recorded and still functioning, though the particular preceding movement itself is past and over" (ibid, 198).

Bringing this all together, Bartlett uses the now famous example of playing tennis. How one swings the racket towards the ball will depend on a large number of things: current visual input and one's current posture, for example, both of which will be influenced by the input/posture that immediately preceded it. Where you end up just before the swing is a result of all the inputs/movements leading up to it. As such, when

you make the swing, you are not simply reproducing a previous set or pattern of behavior exactly as you always do – each instance of performance is unique and “literally manufactured out of the living visual and postural ‘schemata’ of the moment and their interlocutions” (ibid, 202).

In short, remembering will involve putting together schemata and intentions across the background of movement in new ways in every instance. I think Bartlett’s account has several advantages that are worth noting. First, the schemata, or adaptive responses based on past experience ‘en masse’, seem to be understood as basic and importantly based around bodily movement. As such, they are intricately tied to the subject’s past experience, interests, and whatever is going on around them in a way that helps give us both the sense of a ‘rolling mean’ of experience we have been highlighting, and how this mean must be responsive to everything around it. This also gives Bartlett’s account a sense of how procedural memory would include other, likely declarative, elements in its groupings and relation to other elements of past experience.

But the important distinction, here, is that the ‘mean’ or ‘schema’ has an important basis in bodily movement, rather than a broader sense of ‘past experience’. This opens the door to a clearer understanding of the contribution of the physical body to skill – the movement of the body, rather than just its position in a circumstance, helps determine the schema. This foundation in movement also facilitates an account of temporality that draws movements together. Rather than understanding movement as a serial sequence of positions in the ‘right order’, the experience of temporality is based in the body moving in relation to its previous positions. Using previous positions as a

‘background’, one can keep track of other movements and give guidance to the rest of an unfolding schema.

Even so, Bartlett’s account still has some worries similar to those brought up for the enactivist account. For instance, though it does offer a physical grounding in movement that the enactivist account lacks, this grounding is glossed over and not nearly as developed as other accounts we shall soon see. Similarly, though temporality seems to be grasped relative to the moving body, one might worry that Bartlett’s account gives a sense of how temporality is linked *within* a movement but an incomplete sense of its workings *between* movements. Though I think his account has the means to account for this, his own writing on this topic is not as developed as we might like and would require some speculation on our part (especially keeping in mind that this reading of Bartlett is not the typical one). Still, all of this gives us another embodied account that, I think points us in the right direction and helps set the stage for our ultimate exploration.

Chapter 7:

Body and Time – An account of procedural memory

7.0: Introduction - Memory and Time:

The relationship between memory and time has long been noted but how this relationship gets fleshed out varies widely depending on the several factors. On more ‘traditional’ readings, especially within the analytic tradition, memory has been largely understood within the preservationist S&R models, and its relationship to time has been one where at least part of its role is to accurately mark and preserve the passage of events in time (usually understood in some linear fashion). In this sense, memory is largely understood as the basis of our understanding of time.

On one prominent line of argumentation, episodic memory has been thought to be the link to an understanding of time insofar as it creates and maintains the notion of something being in the past. In other words, the thought is that, for us to have a sense of something being in the past, we must have episodic memory that directly and immediately marks experiences as having occurred in the past. On this view, we only have a sense of time and its passage because we have memories that help us conceive of there being such a thing (see for instance, Russell, 2009). However, more recent views have pushed against this notion, suggesting that episodic memory, alone, is not enough to account for our understanding of time and that this likely relies on several other functions of memory (see Hoerl, 2017).

Another leading analytic notion of the relationship between memory and time is that memory provides us with a reliable link to, and therefore knowledge of, the past via a

causal connection. That is, the causal order between events (e.g. one event being the cause of another and therefore preceding it) grounds a sense of time, itself independent of human action. Our memories and their connection to the past are, in turn, constituted by causal links “between the memory and the original experience that led to the memory, and between the experience and what the experience was of” (Poidevin, 2017, 225). In so far as these experiences, and their grip on temporality, are reliable, then our memories are also reliable and will provide us with a sense of temporality.

However, in continental traditions, especially phenomenology, the relationship between time and memory is usually understood quite differently. On the whole, phenomenological accounts take memory to be a function that straddles past, present and future: one that brings all of one’s past experience to bear on the present which, itself, is constantly intentionally aimed towards the future. On these accounts, rather than understanding memory as what gives us a grip on temporality, the relationship between our sense of time and the capacity for memory is understood to be deeply interwoven and interdependent rather than one-directional. An upshot of this is that these views tend to reject the understanding of time as a linear series of ‘nows’ or ‘presents’ tied together through simple chronology. Indeed, in criticizing the analytic approach, Merleau-Ponty (Merleau-Ponty, 2012) says that ‘transporting the time of objects’, based in causality and successive ‘nows’, to ourselves simply assumes that the past is present to consciousness (usually via memory) and does not explain how this might work.. If we base an understanding of the past purely in the ‘traces’ (be they physiological or psychical) of memory, one would be presupposing the notion of recognition which, is not only what

these views are meant to explain, but which also can only be gained through already having some familiarity with the past. Instead, he says that analyzing time grants us access to the concrete structure of subjectivity (ibid, 433) and should, in turn, give us a clearer sense of how memory can be understood to bring all the dimensions of time together.

Given everything that we've seen up till now, especially our analysis of procedural memory extending between past, present, and future, I will put the analytic notions aside in favor of the phenomenological understandings. Of course, there is a fair amount of variation amongst these views, too, and so, it is worth exploring some key phenomenological contributions on the matter of time and memory.

7.1 – Phenomenological accounts of time and memory:

7.1.1. – Bergson on Memory:

Starting with Henri Bergson, who is often credited with noting the distinction between what would be called procedural and episodic memory, is a good root for much of what follows (Bergson, 2014). In general, Bergson argues that memory is, fundamentally, the continuous and automatic synthesis of the past and an ever-changing present into a unified whole. This *contraction* memory accumulates just as, using his example, snow accumulates on a snowball as it rolls down the hill. All of this accumulated past is then constantly 'rolling' into one's present and bringing all of one's past experience into it. But this should not be understood as a present being constantly pushed forward, but rather a continuous synthesis of the past with the present – a unified cohesion of the two where they are inseparable from each other. Indeed, this *contraction*

memory is what allows for a sense of duration and for a consciousness that is given as whole and continuous (for a more expansive summary, see Perri, 2017).

This general idea is meant to hold across all the many and varied forms of memory that Bergson identifies. However, in his original distinction between habit memory and recollection (or what we now more commonly call procedural and episodic memory, respectively), Bergson is often interpreted as suggesting that habit memory is not *really* memory since it is not in the business of recollecting or representing past events, like episodic memory (which he calls memory *par excellence*). Nevertheless, there are those, such as Perri (2017), who suggest that, despite there being clear evidence of Bergson thinking the above, he also, at different times, suggests that habit memory is just an extreme version of the contraction previously described. Indeed, Bergson says that habit memory and recollection should be understood as different, but continuous, ‘tones of mental life’ and that, as such, there is a reason to think of habit memory as memory proper – a memory that he refers to as a ‘memory of the body’. For our purposes, then, it seems reasonable to interpret Bergson as offering an account of the relationship between time and procedural memory whereby past experience is contracted in procedural memory and actively part of one’s present - that these memories are not ‘stored’ anywhere in the mind but are active ‘processes’ and ‘movements’.

7.1.2. – Husserl on Time:

Moving on to the notion of time itself, one of the most influential phenomenological descriptions is that of Edmund Husserl (Husserl, 1991). Husserl aims to give an account of our phenomenological sense of inner-time consciousness, or a sense

of how consciousness experiences time. He claims that perception of ‘now’ time-slices would not be enough to grant us a sense of temporally extended objects or of duration. That kind of experience would amount to “a series of unconnected points of experiencing, like a line of pearls” (Gallagher & Zahavi, 2021, p. 79). This being the case, our consciousness must be able to extend beyond the present moment to be co-conscious of moments just past and those about to come. However, Husserl is clear that he means to differentiate this from memory (though he doesn’t explicitly say so, I believe he means episodic memory). That is, he understands memory to be effortful, something one does actively and that this has its own phenomenological description. Our experience of time, however, is not active in this way – though one may need to deliberately recall experiences in the more distant past, one does not have to deliberately retain the immediately preceding moment(s) when one is, for instance, listening to a particular melody (Cerbone, 2006).

In order to explain the nature of this co-consciousness across the past and present, Husserl introduces the notions of ‘retention’ and ‘protention’. Retention is what provides us with a grip on the moments that have just-passed. Rather than retaining the preceding moment as still active in present consciousness or present in the same way as the new ‘now’, consciousness retains a sense of it having just passed or as an experience fading into the past. So, using Husserl’s own example, as one listens to a melody, each passing note is experienced as a ‘now’ moment following the previous one. As each new ‘now’ arrives, the previous one slips into the past, but is still active in the experience of the present moment - not as experientially present but as having just-passed. In this way,

consciousness retains what has just-passed as part of one's intentional structure, rather than as actively present.

A 'protention' on the other hand, is an expectation or anticipation, an indefinite projection of intention, towards whatever is just about to come. In the same example of listening to a melody, in each present 'now' one is intending oneself into the future with expectations of what is about to happen. Indeed, if one's expectations are not met, one may experience a discontinuity or disruption to one's experience – for instance, if one knows the melody one is listening to and the incorrect note is struck, this will likely be noted in some way and impact one's experience. In just about every activity or action one engages in, one usually has a sense of what is about to happen, what is expected, where things are headed. In a manner similar to what we discussed in the case of predictive processing, this kind of anticipation, especially when it is not met, suggests an intentional projection into the immediate future (Cerbone, 2006; Gallagher & Zahavi, 2021). As Merleau-Ponty summarizes it, Husserl's view is one where "time is not a line, but rather a network of intentionalities" (Merleau-Ponty, 2012, p. 440).

These retentions and protentions indicate to Husserl that there is, what he calls, a 'horizontal' structure of experience. That is, in the case of our melody, for instance, whichever note is being currently experienced will be intentionally linked to those experienced as just-passed and those about-to-be-experienced, and these linked experiences constitute a 'horizon' of the currently experienced moment. In other words, the horizon could be understood as the label for how "the current moment of experience 'points to' those further notes as retained or expected" (Cerbone, 2006, p. 27). The

horizons of each one of those currently experienced moments are, according to Husserl, ‘synthesized’ such that all the moments of experience are understood to be about, or part of, the experience of a melody, rather than a series of disconnected notes. This temporal synthesis is, according to Husserl, the foundation of perceptual synthesis and so is a precondition for all consciousness. He calls this internal structure one’s inner time-consciousness because it is completely central the inner-workings of any action (Cerbone, 2006; Gallagher & Zahavi, 2021).

7.2. – Merleau-Ponty on Time and Memory:

7.2.1 – Time:

Merleau-Ponty’s views on time and memory borrow and build from both Bergson’s and Husserl’s views on the same subjects. But before getting into those details, it is worth giving a primer and his larger philosophical goals and project. Merleau-Ponty is situated in a conversation around how to understand the nature of the self, consciousness, and perception and aims to demonstrate the insufficiency of both ‘empirical’ views and ‘intellectualist’ views. Empiricists, for Merleau-Ponty, understand the body and all its doings as part of an analytical system where perception comes down to impersonal stimulus-response mechanisms that are all, to some extent, certain and determinable: if you can understand the mechanisms of stimulus and response, you can understand all behavior. This means they end up reducing the body to an object which, in turn, leaves no room for (perceptual) consciousness. Furthermore, Merleau-Ponty argues that the empiricists make the mistake of thinking it is possible to ‘step outside’ of our perceptual experience and observe it objectively ‘from the outside’, which means the

empiricists are performing an investigation that both depends on their perception and experience and rejects the impact of the very same perceptions and experiences on their inquiries. Intellectualists, on the other hand, reduce all experience to being the construction of a consciousness. Though they do not treat the body or the subject as an object, Merleau-Ponty thinks that their view strips the world and the environment of any real impact on the subject and creates another sort of certain and determinable world: if you can understand the constructive processes of consciousness, you can then understand everything.

Instead, Merleau-Ponty means to provide a third path which can take on board what each of the previous views contribute, while offering a fundamentally different approach that avoids their key mistakes. What he wants is to ‘return’ to a direct and primitive contact we have with the world; one that is prior to, and provides the basis for, all the conceptualizations we make of it in reflection and analysis. Rather than aiming to explain our experience ‘from the outside’, Merleau-Ponty takes on the phenomenological project of trying to understanding it from the ‘inside’, and our pre-objective experience as the key and central target of investigation. Once we make this ‘return’ to pre-objective experience, we will find that our experience is rife with ambiguities, irregularities, and indeterminacies. In short, our experience, and reflection on it, is always perspectival, situated, sedimented – it will always depend on our being in a particular space and time, in a particular body with a particular point of view on the world. If we can focus on our experience as such, we will find that it will help reveal the fundamental structure of our

experience and, centrally to our purposes, the role time and temporality plays in our experience.

Merleau-Ponty's account of time⁶¹ is aimed at understanding the time of a phenomenal body – a body that both already finds itself in a temporal process it does not create but also plays an important role in the temporalizing of its own experience⁶². In one of the lesser read sections of his *Phenomenology of Perception*, Merleau-Ponty lays out his own view of time and begins by building on, and clarifying, Husserl's notion of time. In response to a potential worry that Husserl's view entails another kind of trace-based view (i.e. that each retention that carries through time offers a 'trace' to past experience), Merleau-Ponty suggests that this is not how we should understand Husserl. Indeed, he reads Husserl's point to be an explanation of how we can still hold the past 'in hand' without having to posit it: that is, how we can draw on the past without having to use traces to get through it. When we do reach back into our past to explicitly remember something, this only works because we, in some sense, directly relive the memory and tie it, through the series of successive horizons, to our present moment. To explain this, he uses the example of seeing a rock at the bottom of a clear pond: though my view of the rock goes through all the 'layers' and depth of the water, I can clearly see the rock and

⁶¹ I am greatly indebted to the works of (Camele, 1975; Morris & Maclaren, 2015; Muldoon, 2006) for their aid in helping piece together an understanding of MP's view, though I have diverged from these sources and, I believe, offered a rather different interpretation of his view.

⁶² Usually, since this section is situated within his larger phenomenological project, most interpreters believe that Merleau-Ponty's account of time in this section of the *Phenomenology* is meant to be a description of phenomenological time, rather than him making any metaphysical or ontological claims: that is, that he is making claims about our consciousness or experience of time, rather than about time itself. Interpreters also tend to agree that his later works shift their focus toward making ontological, rather than phenomenological claims. However, there are those that argue that, even at this stage, Merleau-Ponty was making ontological claims and that they should be treated as such (M. R. Kelly, 2015).

have direct access to it. The same would be true of access to our retained past – that is, I can ‘see’ my past directly and feel it as part of my present, but at a distance. The relation of the past to my present is direct but remote, like the rock at the bottom of the pond. As such, Merleau-Ponty reads Husserl as suggesting that my access to moments passed (via their retentions) is a direct and ‘irrecusable acquisition’ (Merleau-Ponty, 2012, p. 442).

As such, Merleau-Ponty interprets Husserl’s view as saying the following: that “there is no multiplicity of connected phenomena [in time], but rather a single phenomenon of flowing” and that “time is the unique movement that harmonizes with itself in all of its parts, just as a gesture envelops all of the muscular contractions that are necessary for producing it” (Merleau-Ponty, 2012, 442). That is, Merleau-Ponty understands Husserl as arguing for a view of time as a cohesive and complete movement, one that is only understood in its entirety rather than through its pieces. As such, he argues that Husserl’s synthesis of retentions should not be understood as a synthesis of identification (or a linking of all the retentions of the same instant to a particular moment in time) but should, instead, be a synthesis of *transition*.

This is what allows Merleau-Ponty to talk about time (and actions) as being complete in themselves. Through the transition synthesis, all the parts of an event or action are held together and, in a sense, are complete from the time they start. Consider an example Merleau-Ponty uses (which we will come back to, shortly): in *Swann’s Way* (Proust, 2003), Proust describes the love that his character Swann has for Odette. Swann enjoys Odette’s presence, her smile, the way she talks, etc. and, at least at first, the love seems ‘pure’ and uncomplicated. However, as the relationship progresses, Swann

becomes extremely possessive and jealous and begins to taint their relationship. Merleau-Ponty argues that, rather than events acting externally on their relationship such as to create or solicit Swann's jealousy, his love "is already jealous, and has been since its beginning" (ibid, 449). In other words, from the moment the relationship started, the roots of the jealousy were present and were constitutive of the manner in which Swann loved Odette – namely, in his attachment and enjoyment of her presence being dependent on him being the only one to do so. As Merleau-Ponty puts it, "there is a certain manner of loving in which the entire destiny of this love can be instantly seen" and that is present from the time their bond is formed, rather than being gradually developed over time (ibid). In this way, the 'future' of the relationship is present in its 'past' – the ending is present in the beginning and, in this sense, the entire relationship is present in any given instant of its history.

This synthesis of transition is something that Merleau-Ponty thinks is correct about Husserl's view and will be something he aims to maintain in his own. However, Merleau-Ponty begins to note a few ways in which he thinks Husserl's view is incomplete or in need of clarification. For one thing, Husserl's phenomenological view suggests that this synthesis is something achieved by the consciousness of the subject. But Merleau-Ponty does not think that this is the case and is rather something we do not have a hand in. Indeed, he says that "it is clear that I am not the author of time, any more than I am the author of my own heartbeats, nor am I the one who takes the initiative of temporalization" (ibid, 451). Once you are born, you are a situated being through whom time flows and this transition is experienced like "a moving milieu that recedes from us,

like the landscape from the window of a train” (ibid, 443). In this sense, time is given to us and is not something we create. This being the case, we are ‘passive’ in at least part of our experience of time.

Even so, this passage is not something that we merely observe from the ‘outside’: our subjectivity effects it by reaching through it towards a past and a future. That is, rather than merely observing a unified flow of time, through our protentions and retentions, we actively cut across time and draw the future and the past into our present. We connect all these dimensions through our actions and, in this way, draw the disparate dimensions of time into our experience without needing to stand ‘outside’ of time to do so. In a sense, the transition synthesis of time is the background against which we can act: the ‘background of time’ will always be receding from us, like the landscape from the window of the train, but we can also cut across this ever-receding background and draw various aspects of it together in our action while still acting against the background in question.

Because we cut across time in action to draw together the past and future into the present, Merleau-Ponty also thinks that this is what allows for memory distortions. As we move through life and through our ever-shifting present, our grip and relation to the past and future is also always shifting. Though, for instance, I may still have direct access to a past retention as I have direct sight of the rock at the bottom of the clear pond, the significance of the retention and how it appears in my present may change over time. If, for example, I start looking at the pond and notice a very intriguing rock at the bottom, I may, more a while, be quite taken with it and have a clear sense of where it is, what it

looks like, etc. But as I get bored with it and maybe look away for a long while, my relationship with it changes and I may find it harder to spot again or to notice its distinctive features. In this way, one's past (or more specifically, its significance and influence) may be ever-changing and in flux, despite being continuously related to one's present: the relation is not static. The same would be true of protentions and the way one notices or seeks out what is coming. In this way, the dimensions of time are not stable in the way that Husserl's view might suggest.

These considerations lead Merleau-Ponty to suggest that the present is privileged in our temporal experience since it is the point from which all protentions and retentions will occur and the intersection of our bodies with space and time. Even so, we should be careful not to understand this privileging of the present as consciousness being the mere observer of a transition from one present to the next. Indeed, Merleau-Ponty says that "I am just as much directed toward this morning or toward the night that is about to arrive, and although my present is surely this present instant, it is also just as much today, this year, or even my entire life" (ibid, 444). That is, our 'direction' through time should not be understood as being 'locked' in the present or as merely being the observer of a movement from one present to the next. Instead, our present should be understood as transcending itself to reach out towards our past and future and draw them together into an action, a goal, or even a life.

For instance, coming back to the example of Swann's love for Odette, we already established that there is a sense in which the whole relationship and its outcomes are present at the moment of its forming. However, in order for this to be true, we must not

think of this as the collapse of all of the distinct moments of time into each other as this would amount to there not being any sense of time at all. There must be distinct moments that are linked together in a continuous movement. This means that it can still be true that Swann gradually comes to learn of himself and the manner of his relationship with Odette through reflecting on the unfolding of distinct events in time. That is, in the moments in which he comes to understand or ‘recognize’ himself in his action, he must be able to get some manner of ‘external’ vantage point on himself and piece together his grip on his past, present and future. His actions, in this case, of reflection, must then draw together these elements into the present for him to be able to ‘see’ himself, and, in a sense, cut across time to do so. Merleau-Ponty, however, is careful to note that this need not suggest that Swann is either ‘in’ time or ‘out’ of it – in either case, it is unclear how he could get a grip on himself. We must say, then, that Swann just *is*, in some sense, time.

Pulling all of this together, we get a sense of the ambiguity that underlies Merleau-Ponty’s understanding of time. That is, that temporality is neither totally given to us nor constructed by us, but both. This being the case, Merleau-Ponty seems to differentiate himself from Husserl by suggesting that, though Husserl is (partially) right, he only captures one side of this equation – the sense in which it is given to us (and even that is not true in exactly the manner Husserl describes it). Since Husserl only focuses on time as a unified flow and transition synthesis, he misses out on the sense in which the dimensions and elements of time are disparate and need to be brought together by a subjectivity in its experience. Similarly, we can see that temporality is crucial for our subjectivity since, without it as the background to our action, we would not be able to act

and get a view on ourselves. In this way, then, we both passively experience the synthesis and actively piece it together – we are both *in* and *out* of time, we are both *active* and *passive* in its being, we are both *not* time and *are* time. And, furthermore, we couldn't be any of these without being all of them at the same time. That is, we can only hold subjectivity together against the passive background of synthesized time, but in holding it together we are also constituting the background against which we are acting. This ambiguity is central to our primordial, or pre-reflective, experience and shows us the true interrelation between subjectivity and temporality.

On these grounds, Merleau-Ponty argues that both time and subjectivity should be understood as having a fundamentally similar structure which incorporate both active and passive elements at all times. Indeed, at times, he says that “time must be understood as a subject, and the subject must be understood as time” (ibid 445). Just as subjects both passively experience time and actively unify parts of it, time, itself, also ‘passively’ flows and ‘actively’ unifies itself. That is, he suggests that the transition synthesis of time can only be understood to happen on the grounds that time transcends itself and unifies itself into the flow in question – it is both the thing acted up and the thing acting, in the same way that a subject is. This is often read as a statement about temporality as a dimension of subjectivity that helps us get a grip on how it can piece itself together (see Camele, 1975, Morris and Maclaren, 2015). But perhaps it is more pertinent to say that for Merleau-Ponty, we are part of time, itself – neither beings that are completely external to it nor beings that completely constitute it, but rather the subject that “takes up or lives time and merges with the cohesion of life” (ibid, 446).

7.2.2 - Memory:

Throughout his work, whenever Merleau-Ponty mentions memory, he is usually (implicitly) referring to (explicit) episodic memory. Within his larger framework, whatever might be attributable to procedural memory seems largely to be a function of his understanding of the body and the subject as drawing itself together through time and its own actions. Even so, for both episodic and procedural memory, the body is importantly involved and helps ground the means of memory in general. However, the notion of memory at play in procedural memory is one closely related to physical action, as the discussion above helps portray and, as such, brings memory into direct contact with his temporal views.

On Merleau-Ponty's reading of episodic memory, it is generally understood to be a retention. To remember something (as to imagine something) requires a connection to, and a consciousness of, the present – the act of remembering is brought into one's present by an inner perception of recollection (or imagination). This means that the unfolding of the past should not be understood as a representation or construction of a past event, but rather a re-living or a re-present-ing (e.g. making the past present) of the past. Since most remembering will be of things or events 'at a distance' from us, with layers of time between then and now, often we must find a way to traverse these layers. Usually, this means we must actualize past events 'little by little and step by step', or find a means to "unfold" a given duration according to its own 'tempo' (ibid, 446). Regardless, as far as

episodic memory goes, no relation to the future is necessary for the re-living of the past (Vallier, 2015)⁶³.

However, the case of procedural memory seems quite different, as it not only relies on the past, but on a dynamic relationship to the present and future. Like Bergson, Merleau-Ponty seems to highlight that there is an important element of *compression* in so far as it is true that the past must be understood as being rolled into the present. This much should also be familiar in the sense of a rolling-mean, discussed up till now, as a means of maintaining past experience without needing to recall individual instances of past experience. But, as we saw in the case of simulation theories and both Husserl's and Merleau-Ponty's views, a sense of a connection to the future, or protentions, seems necessary for the ability to act fluidly and explain a unified sense of action in time. And this is precisely what Merleau-Ponty adds to the Bergsonian account.

In general, we can understand Merleau-Ponty's account of procedural memory as lining up very neatly with his sense of action in temporality in general. Procedural memory, on this view, would be both active and passive, both stable and unstable. In the same way that subjectivity both experiences time as flowing and then steps 'outside of it' to hold together this same passage, we both passively come to experience memory and actively hold it together. We passively experience the transition synthesis of our future to present to past and all our memories therein. One consequence of this is that our past is brought to bear on our present. But we are also active in memory in so far as we hold all

⁶³ This seems to be true of Merleau-Ponty's position in the *Phenomenology*, but, as Vallier (2015) demonstrates, this changed over the course of his career and seems, by in large, to resemble the case of procedural memory much more closely.

of its pieces together through action. That is, we are constantly determining the relevance and perspective on our memories based on our current circumstances and therefore influencing the process of holding memories together. One can only hold procedural memory together against the background of passive memory and is, in that sense, constituting the background of passive memory at the same time.

However, we must add to this account the more specific ways in which the body, and our knowledge of it, plays into the general framework above and helps clarify the notion of procedural memory. Within Merleau-Ponty's framework are a few key notions for our purposes, namely a 'body schema' and a 'habit body'. The 'body schema' according to Merleau-Ponty, is the 'hold' on one's body and the position of each of one's limbs that is an 'indivisible possession' of a subject' (Merleau-Ponty, 2012, 100-101). He also says that the body schema "designates the installation of the first coordinates, the anchoring of the active body in any object, and the situation of the body confronted with its tasks" (ibid, 103). In short, the body schema is an internal sense and experience of one's body in the world, related to other things and one's environment, against which one always has a point of reference for one's own movement and the presence and actions of anything beyond oneself.

The 'habit body'⁶⁴, on the other hand, is understood to be the general accumulation of experience in the body that grounds its standard or unreflective behavior. As one makes their way through the world, they gather or accumulate a number of

⁶⁴ Though this is the predominant reading of a habit-body in Merleau-Ponty's framework, there may be several other notions that capture what he is after or provide means of improving on his notion – see (Behnke, 2015).

repeated patterns, behaviors, comportments, movements, etc. which fit in and around their body schema. The more well integrated they are into one's life, the more 'sedimented' they are and help to inform one's present 'base-line', as it were. This, of course, can change, and need not point to any deep sense of self or ability or the like – it is merely an element of the state of one's body in any given moment and which can, with time, be shifted in any number of directions.

Both of these notions are central to the case of understanding Merleau-Ponty's sense of procedural memory. The habit body is, in a sense, our accumulated past in our present – the means through which the past can be brought to bear on the present without any manner of representation or storage and retrieval. The body schema, on the other hand, incorporates this habit body and is actively directing our present, while simultaneously opening us up to the future. Procedural memory, then, could be understood as a passive synthesis of one's habit body into their body schema, and the active organization of the body schema through action and interaction in an environment and openness to the future.

To make this a bit more concrete, consider our steadfast example of Rapinoe and her bent shots on goal. As she lines up for a set penalty shot, Rapinoe is, in every moment, experiencing the passive flow of time which simultaneously links her to the future and brings her past memories, including her procedural memories, to bear in the present moment. Different tasks and different environments will play on her body schema differently which will influence which memories or past experiences are obscured or revealed in the current situation. In this way, some reactions and patterns of behavior may

be ‘passively’ solicited and rely primarily on her habit body and its trained awareness. However, as her kick unfolds, through her action and intention she is also simultaneously responsible for holding together all the elements of her memory in her action and directing their influence. She is not merely passively responsive to the situation but is also, through her body schema and its openness to the future, piecing and holding together her habit body in ways that are completely dependent on the context and open up new possibilities for behavior. As such, she has constant awareness of her body, can bring new information to bear, and integrate it with her existing ‘habits’.

This is an importantly different sense of retention from what we saw in previous chapters. For one thing, this interrelation between the habit body and the body schema and the way it brings past experience into the present does not require any sense of memories being represented, stored or retrieved in either the habit body or the body schema – the past and future can be active in present behavior through the combination of the passive flow of time and the active intentionality holding it together, which are both cascading through the subject and their body. As such, this view also avoids all talk of a distinction between declarative and non-declarative memory, declarative and procedural knowledge. Because S&R and reconstruction views (and even some embodied views) were based largely around accounting for declarative memory, their accounts had to artificially make space for procedural memory or reduce it to declarative memory. However, on this account, declarative and procedural are indissolubly linked such that it does not make sense to separate them – both are understood to be inextricably at play.

Putting all of this together, we can see how this kind of view can account for all the advantages discerned from those considered in previous chapters. Considering the first section, we can see how the ‘passivity’ involved in procedural memory on this account incorporates the ideas of many of the non-intellectualist views, while the ‘active’ elements bring in the intellectualist leanings. We can even see how this is a hybrid view of sorts – both must be true and necessarily rely on each other for procedural memory to work properly. Indeed, as the AIR model suggested, we turned our attention towards a third element (rather than automaticity of cognitive control) to help explain this hybridity – but rather than focusing on attention, as they do, our current hybrid notion relies on understanding the nature of time and temporality in our subjectivity to do so.

Similarly, this kind of framework can account for the benefits of both the classical reconstructionist and the simulationist views. For instance, as already noted, Merleau-Ponty takes his understanding of time and its relation to our subjectivity to be able to account for the kind of memory distortions that were central to Bartlett’s understanding of reconstructive memory via the changing relation to one’s retentions in different circumstances. Indeed, Merleau-Ponty’s account offers a way of understanding how events are ‘reconstructed’ in the present without having to rely on any manner of preservationist S&R tendencies. As we saw above, the sense in which one ‘reconstructs’ one’s past relative to one’s current circumstance can be seen in the ways the past does not remain idle or static for me – it is constantly being interpreted and reinterpreted based on whatever is happening in the present and my current protentions. Similarly, via protentions, Merleau-Ponty’s account is able to maintain the kind of future-orientation

and ability to foresee ways for acting found in the simulationist readings and also does so without needing to rely on strict causationality of past events to memory. But, again, he does so in a manner that does not reduce procedural or bodily memory to mere motor- or perceptual-acuity (though, of course, these are involved).

Lastly, Merleau-Ponty's view is an improvement on the contemporary embodied views on memory we discussed in so far as it, again, clearly avoids any preservationist S&R understandings and it gives us a much clearer picture on the relationship between the body, intentionality, and time in the process of memory. Similarly, it incorporates the advantages of a rolling mean from Bartlett's understanding of a procedural memory and integrates the notion of a body schema directly into it. Indeed, Merleau-Ponty's view situates the body-schema much more completely in the larger context of subjectivity, intentionality and temporality, while maintaining all the intuitive benefits and general structure of Bartlett's understanding.

Even with all of these advantages, Merleau-Ponty's general framework might strike some critics as still being relatively broad and abstract for our purposes. In particular, that the view of procedural memory in question is not yet sufficiently grounded in the body and its movements and relies too much on the proposed relationship between time and subjectivity (see, for instance, Casey, 2000 and Sheets-Johnstone below). Another way of putting this might be to suggest that even if Merleau-Ponty's account can be understood to give us a sense of the structure of memory of procedural memory, or *how* we remember, it doesn't tell us much about its content or *what* we remember. As such, it might still be thought to suffer a similar fate to those of enactivist

views discussed in the previous chapter. So, in order to preempt and subvert this kind of worry for an ultimate understanding of procedural memory, I suggest we supplement Merleau-Ponty's account of time and memory with a more concrete understanding of bodily memory presented by Maxine Sheets-Johnstone. Though I believe his view has the means to account for much of what she discusses, Merleau-Ponty does not flesh this side of things out as concretely as Sheets-Johnstone and so, regardless of our view of time, her view could be read as a valuable supplement.

7.3 – Maxine Sheets-Johnstone on Kinesthetic Memory

To bring a concrete and focused attention onto the body in a larger conversation around 'bodily memory', Sheets-Johnstone focuses on kinesthetic memory and the kinetic dynamics of a body. In her discussion of kinesthetic memory (Sheets-Johnstone, 2003), Sheets-Johnstone suggests that kinetic dynamics, or learned patterns of action which are kinesthetically recognizable to the subject, are the essence of kinesthetic memory. In short, she claims that kinesthetic memory is structured by these dynamic patterns and that, once these patterns are learned and familiar, they are then integrated into kinesthetic memory. Borrowing from Luria (Luria, 1973), she utilizes the notion of 'kinetic melodies' to help demonstrate the detail and complexity of these patterns and everything that goes along with them. I take her notion of kinesthetic memory to be largely aligned with our own target phenomenon of procedural memory since, as we will see, it is understood as a habitually based but intentionally guided responsiveness that seems to line up very neatly with the picture of procedural memory we have been building up till now.

To begin with, Sheets-Johnstone explains that “kinesthetic memories are not vague, abstract kinetic phantoms but are inscribed in the body as specific bodily dynamics, dynamics that, as enacted, are at once familiar and tailored distinctively to the particular situation at hand” (ibid, 74). These actions are ‘complete’ and held together in much the way a melody is (remember our conversation around Husserl) – for this reason, she borrows the term ‘kinetic melody’. She argues that the act of any learned performance should be understood as “mov[ing] through a dynamic series of coordinated movements that is kinesthetically felt *both* as dynamic and as dynamically familiar” (ibid, original emphasis). For instance, she provides an example of writing out one’s name or signature. The act of writing one’s name requires not only a dynamic pattern of movement but also one that will be familiar to the subject. For instance, they will know that they have to quickly change direction after a given letter, or that they will accentuate the arch of another letter. But, as they go through it, it will kinetically feel like all the other times they have done it – that is, it will be familiar. This familiarity helps guide and situate the action. This suggests that both the kinetics and its familiarity are integral. This leads Sheets-Johnstone to suggest that each set of coordinated movements will come with its own distinctive ‘spatio-temporal-energetic’ dynamic whose distinctiveness and familiarity are primarily a question of kinesthetics.

Once a dynamic pattern of movement is initiated, Sheets-Johnstone suggests it will ‘flow on by itself’. That is, once something has been engrained in kinesthetic memory and the pattern is activated or initiated, it will naturally tend towards fulfilling the pattern (ibid, 75). Citing Husserl, she notes that this will happen via a series of

retentions and protentions over time, thus holding together the various movements in a series (ibid). However, she notes that this does not mean that the dynamic pattern is unresponsive to circumstance. In the case of writing one's signature, for instance, the well-established pattern of movement will be influenced by the room left for one's signature on the page, by the flow and/or ergonomics of the implement one is using, by the surface on which one is writing, etc. In all these cases, the pattern must and will be responsive.

Furthermore, she aims to separate out the automaticity of initiating the sequence from the performance of the total act. She states that merely starting the sequence is not sufficient to generate the entire performance and that one cannot simply activate an 'auto-pilot' and work through the sequence unconsciously or with no oversight. Instead, she says that "the movement that flows forth effortlessly in a coherent dynamic does so because we know and remember the flow in a corporeally felt sense: *we kinetically instantiate what we know kinesthetically*. What is automatic is, in effect, kinesthetic memory. The kinetic melody runs off by itself because a familiar dynamics is awakened in kinesthetic memory and generated by it" (ibid, original emphasis).

Sheets-Johnstone goes on to explain that whatever impulse initiates a dynamic series will not merely be some inert trigger but must, instead, be volitional. Indeed, as we suggested in our criticism of enactivists in the previous chapter, Sheets-Johnstone believes that volition and kinesthetic memory are importantly linked throughout performance. For instance, in volitionally initiating an action, one will be aware of one's bodily comportment, environmental setting, interactions with tools and implements

around oneself, etc. One does not usually simply find oneself chopping vegetables for dinner or brushing their teeth – indeed, one does not usually find that one has suddenly stopped chopping vegetables or suddenly stopped brushing their teeth. In both cases the initiation of these dynamic patterns and their oversight is always contextualized and necessarily in-tune with one’s kinesthetic awareness. Furthermore, far from merely following through a dynamic series of movements, one can voluntarily change the series in a myriad of ways and so change the dynamic of the performance.

From there, Sheets-Johnstone goes on to discuss what she understands to be deficiencies in other theories that aim to capture this same phenomenon of kinesthetic memory. She criticizes several key notions, including that of a ‘motor’ skill or capacity for focusing too heavily on the mechanics of the body without incorporating any notion of a subject or agent. Similarly, she criticizes both the notion of ‘body image’ (a set of intentional states that includes perception, conceptual understandings, and emotions) and of ‘body schema’ (a subpersonal, preconscious system of motor capacities, abilities and habits)⁶⁵ for comparable reasons. First, she suggests that the notion of a ‘body image’ both leaves out the role of kinesthetic memory on performance, but only works by implicitly taking it for granted. ‘Body schema’, on the other hand, does aim to take kinesthetic memory and history into account, but has no actual basis in experience and is really just an explanatory tool, a missing hypothetical entity, that is theorized in order to fill a gap in the framework⁶⁶.

⁶⁵ see (Gallagher, 1986; Gallagher & Cole, 1995) for more on this distinction)

⁶⁶ As briefly mentioned at the beginning of this section, I believe MP does have means to address this concern and this reading is uncharitable, especially given his larger framework and works post

She advocates, instead, for a notion and methodology that do not separate between the neurological and the experiential and instead aims to unite the two. When we take up ‘pointillist’ descriptions (of a body as an object moving from point A to point B), it is easy to fall into mechanistic descriptions of motor functions but also, mechanistic description of memory functions of storage, retrieval, short- and long-term store, hard distinctions between declarative and non-declarative memory, etc. But this, she suggests, crucially leaves out the temporal nature of kinesthetic memory. This kind of memory is not just a pointillist system of remembered sensations but remembered ‘spatio-temporal-energetic’ dynamics. Incorporating this temporal spread means that “[kinesthetic] memory is not a memory of positions but of a whole body dynamic, which is based not on bodily sensations - localized, positional happenings - but on the perception of movement. In short, kinesthetic memories are constituted through and through by dynamic, not sensational, sedimentations” (ibid,89).

Taking all this together, the key elements of Sheets-Johnstone’s view that are relevant for us are the notions of kinesthetic memory as a) sedimented patterns of kinetic movement that are unique and familiar, each possessing distinct spatio-temporal-energetic signatures, b) that these patterns aim towards completion but, c) that they are necessarily tied to intentionality and temporality and cannot be understood as either mere auto-pilot or as a succession of completely isolated and distinct movements. As such, kinesthetic

Phenomenology. Regardless, I include it because I think it helps motivate her view and helps demonstrate the need for a the kind of specificity her view offers.

memory relies on the dynamics of the whole body in space and time, and this memory relies on one's integration across all these dimensions.

In several regards, her view lines up nicely with the view we attributed to Merleau-Ponty above – indeed, the broad strokes are very similar. But, for one thing, what Sheets-Johnstone's view suggests is that there may be no need for the notion of a 'body schema' in our analysis and, instead, adds to our previous understanding a more concrete sense of what might be involved in a 'habit body' that can guide performance. That is, the 'habit body' and its usual way of going about actions could be more carefully articulated as the retention of distinct kinetic patterns that are unique and familiar to the subject. Given both the uniqueness and familiarity of a particular pattern, the subject can then utilize the sequence as a baseline guide while also imposing more controlled direction 'through' it in line with the subject's current circumstances.

Similarly, Sheets-Johnstone's view also helps us see how procedural memory might amount to more than just motor- or perceptual-acuity and could be more concretely understood as a form of knowledge in itself. Consider, for instance, the understanding of procedural memory on enactivist accounts, which seems to rely primarily on connectionist notions of neural networks being strengthened through repeated firing. This seems to reduce procedural knowledge down to mere acuities that can be trained and honed over time. But Sheets-Johnstone's account, on the other hand, relies on a deep familiarity with our bodies and their movement, knowledge of the history and potential of our movement via its experience and retention in our bodies. It does not seem that this familiarity can be reduced to mere acuity. Indeed, it seems to imply a kind of overarching

unity across the movement that helps hold it together and allows for its holistic expression that can only be maintained in and through the body. In that sense, then, the kinetic understanding of our bodies is the knowledge of procedural/kinesthetic memory.

But again, overall, her view seems to line up nicely with what Merleau-Ponty, in fact, suggests about the nature of the relationship between memory, intentionality and time. Just as Merleau-Ponty suggests that memory is both passive and active in its relationship with intentionality, Sheets-Johnstone seems to suggest a similar dynamic. That is, she maintains that the flow of kinesthetic memory is automatic but that these flows of movement can be monitored and adjusted via intentionality. At the same time, she seems to suggest that intentionality can only be implemented against the background of knowledge and familiarity provided by kinesthetic memory (e.g. that one must be operating in and through the body to be able to guide its movements). Though she is not at explicit about the relationship of all this through time, she also highlights the key notions of protentions and retentions as a means of holding memory and intention together over the course of an action. In light of all of this, I take it that their views are compatible and that Sheets-Johnstone's view can be seen as a suitable supplement to Merleau-Ponty's lack of specification in the content of procedural memory.

7.4 – Walking through Time and Memory:

When we originally discussed walking all the way back in Chapter 3, we were focused on it primarily as a means of demonstrating that even very basic actions, including walking, should be understood as skills and, as such, fall under the purview of procedural memory. Now, however, we are well placed to use it as an example for

understanding procedural memory at the most basic level of our physical lives and one of the first sets of procedural memories we acquire.

As an infant begins to move its limbs and gain some basic kinetic and kinesthetic familiarity with their body, they will begin to piece together some sequences of actions and movements in relation to their understanding of their body. These may begin as mere wavings and flailings of various limbs, but all of these will eventually contribute to an understanding of the ways in which their limbs can support propulsion and, eventually, the weight of their bodies. For instance, as they gradually acquire the strength to sit upright and support their heads, and as their legs become strong enough to support the rest of their body, they will begin to experiment with scoots, rolls, inch-worming, belly-crawls, army-crawls and the like. As each of these kinetic movements becomes increasingly familiar to the infant, they will become ‘sedimented’ into the infant’s kinesthetic memory and contribute to their growing repertoire of kinetic movements.

This growing repertoire of kinesthetic memory provides the background against which they can start acting intentionally to draw themselves across space. Their intentionality, at this stage, means that infants are using their kinesthetic memory as a base from which to initiate movement and which they can oversee or monitor according to the needs of the moment. For instance, they may be drawn towards a particular toy across the room and, volitionally initiating their kinesthetic memories of, say, crawling, the infant will then be able to ‘play out’ the series of sedimented kinetic movements and oversee their implementation given the shag carpet along the way or the hard floor on the other side. Each movement has its own set of protentions and retentions tied to it,

bringing the infant's past experience directly to bear in the present and also directing them towards the next stage of their progression.

At the same time, however, the infant's intentionality is holding together the 'background' of kinesthetic memory against which they are acting. That is, depending on the details of the exact moment, the infant is highlighting particular protentions and retentions and flexibly utilizing and unifying them in the course of their action. For instance, maybe, as they move across the floor, the infant encounters a particular obstacle that brings new protentions and retentions to bear and that they must unify in order to continue. The pillowing lying in their path may bring to bear new retentions of different kinetic movements that would help them get over the pillow or get under it or scoot around it and will also set new protentions of what to expect as they undergo any one of those possibilities. In this way, their intentionality is unifying past, present, and future in new ways in every instance while also being completely reliant on this background to do so.

As the infant encounters new and unfamiliar circumstances, and as they deploy their existing kinesthetic memories in new ways, they will eventually develop new kinesthetic memories as well. The infant may still have and use older retained kinetic patterns that are solicited in different situations but will likely have different vantage points on them and their usefulness and will also likely develop new sets of protentions and retentions that solicit different kinesthetic memories in different circumstances. In time, through this kind of process, these early forms of locomotion will pave the path for cruising and then for walking. In the end, the processes that led up to walking, and even

walking itself, become so sedimented in kinesthetic memory that they become foundational for a nearly infinite array of other actions one ‘forgets’ the manner the extent to which it was learned or acquired.

7.5 – Conclusion – and the appropriate metaphor?

With all this in hand, let us then summarize our understanding of procedural memory up to this point. We first began this project with an extensive exploration of the nature of skill to determine some criteria for skill that would be the basis of our understanding of what an account of procedural memory, or memory of skill, should be able to encapsulate. From there, we examined several leading notions of (declarative) memory and aimed to determine their application to the procedural case. After having established that it could not be a question of storage and retrieval, or reconstruction or simulation, we did note that, at least the latter two options started pointing towards some promising directions. Along with some notions of how past experience may be ‘rolled in’ to present experience and the ways in which procedural memory may be importantly future-oriented, we suggested that some important basis in the body was required for a more complete sense of procedural memory. However, having examined some contemporary embodied views, we noted that many of their notions of memory either still relied on preservationist S&R understandings or took procedural memory for granted and built frameworks around it, rather than explaining it.

All of this pushed us towards some phenomenological understandings, especially that of Merleau-Ponty, for a grounding of memory both in the body and in time. Building on Bergson and Husserl, we saw how Merleau-Ponty put together an account of

procedural memory as an ambiguously active and passive process of subjectivity through time: a process of passively experiencing the flow of time and, through protentions and retentions, the maintenance of our past in our present and a directedness towards the future. At the same time, memory is, for Merleau-Ponty, an active process of holding together the passage of the flow via our intentionality and action, which also accounts for the many distortions that are rife in memory. Even so, we noted that this account of procedural memory may only describe the structure of this kind of memory and not its content. As such, we brought in the account of kinesthetic memory described by Sheets-Johnstone to suggest that the proper grounding of procedural memory is in kinesthetic knowledge of, and familiarity with, the kinetic sequences of our bodies. These, too, are held together via intentionality through time, and kinesthetic memory and intentionality are integrally intertwined.

As such, our concluding notion of procedural memory is that of a fundamental element of all action where the body, intentionality, and time must interact to enable performance. The knowledge or contents of the memory that is utilized in the process of procedural memory is fundamental knowledge of the kinetics of one's body and its sedimented ways of movement. This kinesthetic memory draws one's body across and through time, both bringing the past into one's present and directing one forward through protentions and retentions. As such, it is also the necessary background upon which intentionality can work, overseeing and monitoring the unfolding of one's actions. At the same time, intentionality is also shaping the 'passive' background of kinesthetic memory and so actively reshaping the ground on which it is based. This whole process

encapsulates everything we've been meaning to capture in an understanding of procedural memory: this whole process is procedural memory.

Considering all of this, is there an appropriate metaphor for discussing our current understanding? Having clearly established the insufficiency of the S&R metaphor, including its development into the reconstructive and simulationist metaphors (at least for our purpose of understanding procedural memory), it is worth discussing what an appropriate metaphor may actually be. There are a number to consider, beginning with, for instance, the 'shaping' metaphor that comes from the literature on behaviorism. On this metaphor, memory is more like a process of shaping than storing, retrieving, reconstructing, or the like. That is, subjects and their memories are shaped over time such that gradual accumulations of past experience constitute new shifts in the item taking shape. For instance, recall the snowball example discussed earlier by Bergson. Similarly, suppose we conceive of a ball of wax being shaped over time as maybe the addition or subtraction of small bits of wax, or the literal reshaping of ball from a ball to a cube, or what have you. When 'pressed' or 'triggered' in its environment, wax's 'reactions' will be based on the manner of its present shaping: if its round, it may roll; if it is square it may turn over a side or two; if it is light it may roll rather far, while if it has become heavy it may not roll far at all, etc. In this way, one's procedural memories may be thought to 'shape' a subject such that, depending on their current state and the nature of their interactions with their environment, they will be disposed to react in the manner set out by the interplay of their shaping and the 'stimulus'.

Though this metaphor does seem to nicely capture the sense of a rolling mean we had been discussing up till now, and to do so without any sense of stored or retrieved information. The ‘shaping’ of the subject is not only ‘stored’ in it in so far as alterations are ‘stored’ (which seems to be a stretch of the notion), and are certainly not retrieved but solicited. The main worry with this metaphor is that it seems to leave out the sense of intentionality and activity on the side of the subject. That is, all the shaping of the wax, say, happens not through the wax’s actions, but presumably at the hands of some other active agent. Of course, we could imagine, for instance, that the ball is not being shaped by an agent, but perhaps a more passive environment – for instance, maybe it fell off its table and is now rolling around the room picking up dust, other bits of wax or dirt on the floor, etc. But in this case too, the impetus for its movement is passive and relies on something else driving it. For this reason, this metaphor seems insufficient for understanding procedural memory which, as we’ve established, is critically intentional and driven by the subject.

There is also a metaphor we considered early on in this project: the notion of the Banyan tree as a model for skill acquisition (Adolph & Robinson, 2013). In this metaphor, the acquisition of skill is likened to the growing process of the banyan tree whereby its roots grow down from the parent branch on which it starts growing, and will depend on many aspects of its environment: that is, as it encounters obstacles it can change its direction of growth to move around, or sometimes even through, them. In time, the original roots become trunks and the later roots and branches grow in all directions, in and out of each other. The unique circumstances of each tree will determine the more

specific growth of all its parts, while also maintaining the general pattern of growth specific to its species.

As an understanding of procedural memory, we might consider this a variety of connectionist neural-network type view. That is, as a subject gradually practices begins the process of acquiring a skill, they also gradually develop procedural memories about this. The roots and branches of their memory will grow in response to whatever they encounter and will be shaped by the intentionality of the subject. Whatever roots or branches get strengthened over time will depend on the position and specific conditions of the subject – their particular history and responsiveness, as well as the particular inclinations determined by their current positioning, will influence which procedural ‘limbs’ get strengthened. As such, this process of memory will be influenced by the subject’s intentionality, their accumulated ‘history’, their environment and reactivity in it, and the like.

This does seem appealing on many levels but, as we suggested in the previous chapter, this kind of view can have enticing properties for those grounded in embodied viewpoints. For instance, the enactivists seemed to endorse a version of this kind of connectionist understanding in so far as it could potentially explain how memories could be reactivated without needing any kind of representation, and we could add to that they might endorse it in so far as this version of it actively includes an account of how this development would include important relations to intentionality and environment. However, though this model might be very helpful for a sense of the acquisition of procedural memory and its instantiation in the brain, it does little to *explain* the

relationships it takes for granted. In particular, it is unclear, on this view, how these factors are meant to come into play in real-time performance in any given instant. As such, it may be useful for one dimension of procedural memory, but not others.

Another view we might consider is that of (Wechsler, 1963), who suggests a rather ‘extreme’ interpretation. In short, he explicitly rejects the storage and search metaphors and instead, endorses a handful of historical views that suggest “[m]emories are not like filed letters stored in cabinets or unhung paintings in the basement of a museum. Rather, they are like melodies realized by striking the keys on a piano. Ideas are no more stored in the brain than melodies in the keys of a piano” (Wechsler, 1963, 150-151). Wechsler adds to this saying that “the brain probably consists of many different keyboards, intercommunicating to be sure, but not always in tune and, of course, not all with a Steinway quality. Some, I am afraid, remind us more of penny whistles or cigar-box ukes” (ibid).

I think there is something hugely appealing about his metaphor for the following reasons: again, it does not rely on any sense of storage or retrieval, or indeed reconstruction or simulation. When played, the piano simply emits the sounds directly and there is no sense in which the sounds are ‘preserved’ in the piano. Instead, on this view, memory is more about potential and actualized potential that can be drawn out in any number of ways, styles, etc. The memories are only ‘there’ whenever the subject is ‘being played’ and, though they can be revived at a later time by playing the same notes, each time will likely be a little different – maybe more or less in tune, maybe with different types or degrees of emotional expression, etc. As we have seen many times, the

notion of melodies being held together through time and intentionality is a hugely appealing one to many thinkers and thinking of memories similarly seems equally so.

However, this metaphor, too, seems to suffer from a particular problem, namely what seems like a mono-directionality in influence. That is, it is obvious how the musician, or subject, can influence the piano, or memories, and bring them to bear. But it is not so obvious how the piano, or memories, influence and lay the ground for the subject to act. Of course, there must be a level of responsiveness to the ‘notes’ and the melody as they are played, but this metaphor, alone, does not give a sense of how the accumulated experience of one’s past can be brought to bear on one’s present and how this will necessarily influence the implementation of one’s intentionality and performance. This is either taken for granted or would need some other metaphor (maybe a sense of how the piano could be ‘shaped’ over time or what have you) to supplement it and help account for this aspect of procedural memory.

And so, it seems the appropriate metaphor for procedural memory has not yet been found and the search continues...

References:

- Abdulrahman, A., & Zebas, C. (1993). Kinematic and Kinetic Comparison of Lean and Obese. *International Symposium on Biomechanics in Sports*, 11, 98–101.
- Acredolo, L. P., Adams, A., & Goodwyn, S. W. (1984). The role of self-produced movement and visual tracking in infant spatial orientation. *Journal of Experimental Child Psychology*, 38(2), 312–327. [https://doi.org/10.1016/0022-0965\(84\)90128-0](https://doi.org/10.1016/0022-0965(84)90128-0)
- Adams, F. (2011). Husker Du. *Philosophical Studies*, 153(1), 81–94. <https://doi.org/doi:10.1007/s11098-010-9663-4>
- Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future : Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, 45, 1363–1377. <https://doi.org/10.1016/j.neuropsychologia.2006.10.016>
- Adolph, K. E., & Robinson, F. R. (2013). The Road to Walking: What Learning to Walk Tells Us about Development. In P. Zelazo (Ed.), *Oxford Handbook of Developmental Psychology*. Oxford University Press.
- Adolph, Karen E. (1995). Psychophysical Assessment of Toddlers' Ability to Cope With Slopes. *Journal of Experimental Psychology: Human Perception and Performance*, 21(4), 734–750.
- Adolph, Karen E. (2000). Specificity of learning: Why Infants Fall over a Veritable Cliff. *Psychological Science*, 11(4), 290–295. <https://doi.org/10.1111/1467-9280.00258>
- Adolph, Karen E., & Avolio, A. M. (2000). Walking infants adapt locomotion to changing body dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, 26(3), 1148–1166. <https://doi.org/10.1037/0096-1523.26.3.1148>
- Adolph, Karen E., Berger, S. E., & Leo, A. J. (2011). Developmental continuity? Crawling, cruising, and walking. *Developmental Science*, 14(2), 306–318. <https://doi.org/10.1111/j.1467-7687.2010.00981.x>
- Adolph, Karen E., Cole, W. G., Komati, M., Garciaguirre, J. S., Badaly, D., Lingeman, J. M., Chan, G. L. Y., & Sotsky, R. B. (2012). How Do You Learn to Walk? Thousands of Steps and Dozens of Falls per Day. *Psychological Science*, 23(11), 1387–1394. <https://doi.org/10.1177/0956797612446346>

- Adolph, Karen E., Joh, A. S., & Eppler, M. A. (2010). Infants' Perception of Affordances of Slopes Under High and Low Friction Conditions. *Journal of Experimental Psychology: Human Perception and Performance*, 36(4), 797–811.
<https://doi.org/10.1037/a0017450>.Infants
- Adolph, Karen E., Karasik, L. B., & Tamis-LeMonda, C. S. (2010). Motor Skill. In M. H. Bornstein (Ed.), *Handbook of Cultural Developmental Science* (pp. 61–88). Taylor & Francis.
- Adolph, Karen E., Tamis-LeMonda, C. S., Ishak, S., Karasik, L. B., & Lobo, S. A. (2008). Locomotor Experience and Use of Social Information Are Posture Specific. *Developmental Psychology*, 44(6), 1705–1714.
<https://doi.org/10.1037/a0013852>.Locomotor
- Adolph, Karen E., Vereijken, B., & Denny, M. A. (1998). Learning to Crawl. *Child Development*, 69(5), 1299–1312. <https://doi.org/10.1111/j.1467-8624.1998.tb06213.x>
- Adolph, Karen E., Vereijken, B., & Shrout, P. E. (2003). What Changes in Infant Walking and Why. *Child Development*, 74(2), 475–497.
<https://doi.org/10.1111/1467-8624.7402011>
- Adolph, Karen E., Bertenthal, B. I., Boker, S. M., Goldfield, E. C., & Gibson, E. J. (1997). Learning in the Development of Infant Locomotion. *Monographs of the Society for Research in Child Development*, 62(3).
- Aldridge, J. W., & Berridge, K. C. (2003). Basal Ganglia Neural Coding of Natural Action Sequences. In A. M. Graybiel, M. R. DeLong, & S. T. Kitai (Eds.), *The basal ganglia VI* (pp. 279–287). Plenum Press. https://doi.org/10.1007/978-1-4615-0179-4_28
- Andre-Thomas, & Autgaerden, S. (1966). *Locomotion from Pre- to Post-natal Life* (Issue 24). The Lavenham Press.
- Assaiante, C. (1998). Development of locomotor balance control in healthy children. *Neuroscience and Biobehavioral Reviews*, 22(4), 527–532.
[https://doi.org/10.1016/S0149-7634\(97\)00040-7](https://doi.org/10.1016/S0149-7634(97)00040-7)
- Audi, R. (2002). The Sources of Knowledge. In P. K. Moser (Ed.), *The Oxford Handbook of Epistemology* (pp. 71–94). Oxford University Press.
<https://doi.org/doi:10.1093/oxfordhb/9780195301700.003.0003>
- Augustine. (n.d.). *Confessions* (A. C. Outler (Ed.)).

- Barnier, A. J., & Sutton, J. (2008). From individual to collective memory : Theoretical and empirical perspectives. *Memory*, 16(3), 177–182.
<https://doi.org/10.1080/09541440701828274>
- Bartlett, F. C. (1995). *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press. <https://doi.org/10.1080/00207284.1951.11507883>
- Bastien, G. J., Schepens, B., Willems, P. A., & Heglund, N. C. (2005). Energetics of Load Carrying in Nepalese Porters. *Science*, 308, 1755.
- Bayley, N. (1969). *Manual for the Bayley Scales of Infant Development*. The Psychological Corporation.
- Bayley, Nancy. (1965). Comparisons of Mental and Motor Test Scores for Ages 1-15 Months by Sex , Birth Order , Race , Geographical Location , and Education of Parents. *1 Child Development*, 36(2), 379–411.
- Behnke, E. A. (2015). Null-Body, Protean Body, Potent Body, Neutral Body, Wild Body. In D. Morris & K. Maclaren (Eds.), *Time, Memory, Institution: Merleau-Ponty's New Ontology of Self* (pp. 69–90). Ohio University Press.
- Beilock, S. L., Bertenthal, B. I., McCoy, A. M., & Carr, T. H. (2004). Haste does not always make waste: Expertise, direction of attention, and speed versus accuracy in performing sensorimotor skills. *Psychonomic Bulletin and Review*, 11(2), 373–379.
<https://doi.org/10.3758/PBR.16.3.463>
- Beilock, S. L., & Carr, T. H. (2004). From novice to expert performance: Memory, attention and the control of complex sensori-motor skills. In *Skill Acquisition in Sport* (pp. 309–327). Routledge.
- Beilock, S. L., Carr, T. H., Macmahon, C., & Starkes, J. L. (2002). When Paying Attention Becomes Counterproductive : Impact of Divided Versus Skill-Focused Attention on Novice and Experienced Performance of Sensorimotor Skills. *Journal of Experimental Psychology: Applied*, 8(1), 6–16. <https://doi.org/10.1037//1076-898X.8.1.6>
- Beilock, S. L., & Gray, R. (2007). Why Do Athletes Choke under Pressure? In G. Tenenbaum & R. C. Eklund (Eds.), *Handbook of Sport Psychology* (3rd ed., pp. 425–444). Wiley.
- Ben-Ze'ev, A. (1992). Cognitive Development: Two Paradigms. In H.-G. Geissler, S. W. Link, & J. T. Townsend (Eds.), *Cognition, Information Processing, and Psychophysics* (pp. 67–90). Lawrence Erlbaum Associates.

- Ben-Ze'ev, A. (1996). The alternative to the storehouse metaphor. *Behavioral and Brain Sciences*, 19(2), 192–193.
- Benson, J. B., & Užgiris, I. Č. (1985). Effect of Self-Initiated Locomotion on Infant Search Activity. *Developmental Psychology*, 21(6), 923–931. <https://doi.org/10.1037/0012-1649.21.6.923>
- Berger, S. E., & Adolph, K. E. (2003). Infants Use Handrails as Tools in a Locomotor Task. *Developmental Psychology*, 39(3), 594–605. <https://doi.org/10.1037/0012-1649.39.3.594>
- Berger, S. E., Adolph, K. E., & Kavookjian, A. E. (2010). Bridging the gap: Solving spatial means-ends relations in a locomotor task. *Child Development*, 81(5), 1367–1375. <https://doi.org/10.1111/j.1467-8624.2010.01478.x>
- Berger, S. E., Adolph, K. E., & Lobo, S. A. (2005). Out of the toolbox: Toddlers differentiate wobbly and wooden handrails. *Child Development*, 76(6), 1294–1307. <https://doi.org/10.1111/j.1467-8624.2005.00851.x>
- Berger, S. E., Theuring, C., & Adolph, K. E. (2007). How and when infants learn to climb stairs. *Infant Behavior and Development*, 30(1), 36–49. <https://doi.org/10.1016/j.infbeh.2006.11.002>
- Bergson, H. (2014). *Henri Bergson: Key Writings* (K. A. Pearson & J. O. Maoilearca (Eds.)). Bloomsbury.
- Bernecker, S. (2008). *The Metaphysics of Memory* (S. Hetherington (Ed.)). Springer.
- Bertenthal, B., Campos, J. J., & Barrett, K. C. (1984). Self-produced Locomotion: an organizer of emotional, cognitive, and social development in infancy. In R. N. Emde (Ed.), *Continuities and Discontinuities in Development* (pp. 175–176). Plenum Press.
- Bertenthal, B. I., Boker, S. M., & Xu, M. (2000). Analysis of the perception-action cycle for visually induced postural sway in 9-month-old sitting infants. *Infant Behavior and Development*, 23(3–4), 299–315. [https://doi.org/10.1016/S0163-6383\(01\)00046-7](https://doi.org/10.1016/S0163-6383(01)00046-7)
- Bertenthal, B. I., & Campos, J. J. (1984). A Reexamination of Fear and Its Determinants on the Visual Cliff. *The Society for Psychophysiological Research*, 21(4), 413–417.

- Bertenthal, B. I., Rose, J. L., & Bai, D. L. (1997). Perception-Action Coupling in the Development of Visual Control of Posture. *Journal of Experimental Psychology: Human Perception and Performance*, 23(6), 1631–1643. <https://doi.org/10.1037/0096-1523.23.6.1631>
- Biringen, Z., Emde, R. N., Campos, J. J., & Appelbaum, M. (2008). Development of Autonomy: role of walking onset and its timing. *Perceptual and Motor Skills*, 106, 395–414.
- Biringen, Z., Emde, R. N., Campos, J. J., & Appelbaum, M. I. (1995). Affective Reorganization in the Infant, the Mother and the Dyad: the Role of Upright Locomotion and Its Timing. *Child Development*, 66, 499–514.
- Boyd, L. A., Edwards, J. D., Siengsukon, C. S., Vidoni, E. D., Wessel, B. D., & Linsdell, M. A. (2009). Motor sequence chunking is impaired by basal ganglia stroke. *Neurobiology of Learning and Memory*, 92(1), 35–44. <https://doi.org/10.1016/j.nlm.2009.02.009>
- Bril, B., & Brenière, Y. (1992). Postural requirements and progression velocity in young walkers. *Journal of Motor Behavior*, 24(1), 105–116. <https://doi.org/10.1080/00222895.1992.9941606>
- Bril, B., & Ledebt, A. (1998). Head coordination as a means to assist sensory integration in learning to walk. *Neuroscience and Biobehavioral Reviews*, 22(4), 555–563. [https://doi.org/10.1016/S0149-7634\(97\)00044-4](https://doi.org/10.1016/S0149-7634(97)00044-4)
- Bril, B., & Sabatier, C. (1986). The Cultural Context of Motor Development: Postural Manipulations in the Daily Life of Bambara Babies (Mali). *International Journal of Behavioral Development*, 9, 439–453.
- Brown, A. D., Kouri, N., & Hirst, W. (2012). Memory's malleability: its role in shaping collective memory and social identity. *Frontiers in Psychology*, 3(JUL), 1–3. <https://doi.org/10.3389/fpsyg.2012.00257>
- Browning, R. C., & Kram, R. (2007). Effects of obesity on the biomechanics of walking at different speeds. *Medicine and Science in Sports and Exercise*, 39(9), 1632–1641. <https://doi.org/10.1249/mss.0b013e318076b54b>
- Browning, R. C., & Kram, R. (2009). Pound for pound: Working out how obesity influences the energetics of walking. *Journal of Applied Physiology*, 106(6), 1755–1756. <https://doi.org/10.1152/jappphysiol.00373.2009>

- Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6(4), 345–375.
<https://doi.org/10.1037/h0073117>
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The Brain's Default Network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124, 1–38. <https://doi.org/10.1196/annals.1440.011>
- Bushnell, E. W., McKenzie, B. E., Lawrence, D. A., & Connell, S. (1995). The Spatial Coding Strategies of One-Year-Old Infants in a Locomotor Search Task. *Child Development*, 66(4), 937–958. <https://doi.org/10.1111/j.1467-8624.1995.tb00914.x>
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and Motor Representation in Purposive Action. *Philosophy and Phenomenological Research*, 88(1), 119–145.
<https://doi.org/10.1111/j.1933-1592.2012.00604.x>
- Byrne, A. (2010). Recollection, Perception, Imagination. *Philosophical Studies*, 148(1), 169–186. <https://doi.org/doi:10.1007/s11098-010-9508-1>
- Cabeza, R., Rao, S. M., Wagner, A. D., Mayer, A. R., & Schacter, D. L. (2001). Can medial temporal lobe regions distinguish true from false ? An event-related functional MRI study of veridical and illusory recognition memory. *Proceedings of the National Academy of Sciences of the United States of America*, 98(8), 4805–4810. <https://doi.org/10.1073/pnas.081082698>
- Camele, A. M. (1975). Time in Merleau-Ponty and Heidegger. *Philosophy Today*, 19(3), 256–268.
- Campos, J. J., & Anderson, D. I. (2000). Travel broadens the mind. *Infancy*, 1(2), 149–219. <https://doi.org/10.2307/j.ctv5rf2m8.6>
- Campos, J. J., Bertenthal, B. I., & Kermoian, R. (1992). Early Experience and Emotional Development: The Emergence of Wariness of Heights. *Psychological Science*, 3(1), 61–64. <https://doi.org/10.1111/j.1467-9280.1992.tb00259.x>
- Campos, J. J., Hiatt, S., Ramsay, D., Henderson, C., & Svejda, M. (1978). The Emergence of Fear on the Visual Cliff. In M. Lewis (Ed.), *The Development of Affect* (pp. 149–150). Springer Science and Business Media.
<https://doi.org/10.4135/9781412972024.n2655>
- Capute, A. J., Shapiro, B. K., Palmer, F. B., Ross, A., & Wachtel, R. C. (1985). Normal Gross Motor Development: the Influences of Race, Sex and Socio-Economic Status. *Developmental Medicine & Child Neurology*, 27(5), 635–643.
<https://doi.org/10.1111/j.1469-8749.1985.tb14136.x>

- Carrier, D. R., Kapoor, A. K., Kimura, T., Nickels, M. K., Scott, E. C., So, J. K., & Trinkaus, E. (1984). The Energetic Paradox of Human Running and Hominid Evolution [and Comments and Reply]. *Current Anthropology*, 25(4), 483–495. <https://doi.org/10.1086/203165>
- Casey, E. S. (2000). *Remembering: A Phenomenological Study* (Second). Indiana University Press.
- Cavagna, G. A., Heglund, N. C., & Taylor, C. R. (1977). Mechanical work in terrestrial locomotion: two basic mechanisms for minimizing energy expenditure. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 233(5), R243–R261.
- Cerbone, D. R. (2006). *Understanding Phenomenology*. Acumen.
- Chang, C. L., Kubo, M., Buzzi, U., & Ulrich, B. (2006). Early changes in muscle activation patterns of toddlers during walking. *Infant Behavior and Development*, 29(2), 175–188. <https://doi.org/10.1016/j.infbeh.2005.10.001>
- Chen, L. C., Metcalfe, J. S., Jeka, J. J., & Clark, J. E. (2007). Two steps forward and one back: Learning to walk affects infants' sitting posture. *Infant Behavior and Development*, 30(1), 16–25. <https://doi.org/10.1016/j.infbeh.2006.07.005>
- Chisholm, R. M. (1989). *Theory of Knowledge* (3rd ed.). Prentice Hall International.
- Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology*, 32(5), 692–718. <https://doi.org/10.1080/09515089.2019.1607279>
- Christensen, W., Sutton, J., & McIlwain, D. J. F. (2016). Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience. *Mind and Language*, 31(1), 37–66. <https://doi.org/10.1111/mila.12094>
- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action and Cognitive Extension*. Oxford University Press.
- Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind* -. Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, January, 7–19.
- Clark, J. E., Whittall, J., & Phillips, S. J. (1988). Human interlimb coordination: The first 6 months of independent walking. *Developmental Psychobiology*, 21(5), 445–456. <https://doi.org/10.1002/dev.420210504>

- Clearfield, M. W. (2004). The role of crawling and walking experience in infant spatial memory. *Journal of Experimental Child Psychology*, 89(3), 214–241. <https://doi.org/10.1016/j.jecp.2004.07.003>
- Clearfield, Melissa W., Osborne, C. N., & Mullen, M. (2008). Learning by looking: Infants' social looking behavior across the transition from crawling to walking. *Journal of Experimental Child Psychology*, 100(4), 297–307. <https://doi.org/10.1016/j.jecp.2008.03.005>
- Corbetta, D., & Bojczyk, K. E. (2002). Infants return to two-handed reaching when they are learning to walk. *Journal of Motor Behavior*, 34(1), 83–95. <https://doi.org/10.1080/00222890209601933>
- Craik, F. I. M., & Lockhart, R. S. (1972). Levels of Processing: A Framework for Memory Research. *Journal of Verbal Learning and Verbal Behavior*, 11, 671–684. <https://doi.org/10.4324/9781315440446>
- Csikszentmihalyi, M. (1990). *Flow: the Psychology of Optimal Experience*. Harper Perennial.
- De Brigard, F. (2014a). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191(2), 155–185. <https://doi.org/10.1007/s11229-013-0247-7>
- De Brigard, F. (2014b). The nature of memory traces. *Philosophy Compass*, 9(6), 402–414. <https://doi.org/10.1111/phc3.12133>
- De Brigard, F., Spreng, R. N., Mitchell, J. P., & Schacter, D. L. (2015). NeuroImage Neural activity associated with self, other, and object-based counterfactual thinking. *NeuroImage*, 109, 12–26. <https://doi.org/10.1016/j.neuroimage.2014.12.075>
- de Vries, J. I. P., Visser, G. H. A., & Prechtl, H. F. R. (1982). The emergence of fetal behaviour. I. Qualitative aspects. *Early Human Development*, 7(4), 301–322. [https://doi.org/10.1016/0378-3782\(82\)90033-0](https://doi.org/10.1016/0378-3782(82)90033-0)
- de Vries, J. I. P., Visser, G. H. A., & Prechtl, H. F. R. (1985). The emergence of fetal behaviour. II. Quantitative aspects. *Early Human Development*, 12(2), 99–120. [https://doi.org/10.1016/0378-3782\(85\)90174-4](https://doi.org/10.1016/0378-3782(85)90174-4)
- Debus, D. (2014). 'Mental Time Travel': Remembering the Past, Imagining the Future, and the Particularity of Events. *Review of Philosophy and Psychology*, 5(3), 333–350. <https://doi.org/10.1007/s13164-014-0182-7>

- Debus, D. (2017). Memory Causation. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory*. Routledge.
- Dickie, I. (2012). Skill Before Knowledge. *Philosophy and Phenomenological Research*, 85(3), 737–745.
- Dijkstra, K., Kaschak, M. P., & Zwaan, R. A. (2007). Body posture facilitates retrieval of autobiographical memories. *Cognition*, 102, 139–149.
<https://doi.org/10.1016/j.cognition.2005.12.009>
- Dominici, N., Daprati, E., Nico, D., Cappellini, G., Ivanenko, Y. P., & Lacquaniti, F. (2009). Changes in the limb kinematics and walking-distance estimation after shank elongation: Evidence for a locomotor body schema? *Journal of Neurophysiology*, 101(3), 1419–1429. <https://doi.org/10.1152/jn.91165.2008>
- Draaisma, D. (2000). *Metaphors of Memory: A History of Ideas About the Mind* (P. Vincent (Ed.)). Cambridge University Press.
- Dreyfus, H., & Dreyfus, S. (1986). *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. The Free Press.
- Dreyfus, H., & Kelly, S. D. (2007). Heterophenomenology: Heavy-handed sleight-of-hand. *Phenomenology and the Cognitive Sciences*, 6(1–2), 45–55.
<https://doi.org/10.1007/s11097-006-9042-y>
- Dreyfus, H. L. (1997). Intuitive, Deliberative, and Calculative Models of Expert Performance. In C. E. Zsombok & G. Klein (Eds.), *Naturalistic Decision Making* (pp. 17–28). Lawrence Erlbaum Associates.
- Dreyfus, H. L. (2002). Intelligence without representation - Merleau-Ponty's critique of mental representation. *Phenomenology and the Cognitive Sciences*, 1(4), 367–383.
<http://philpapers.org/rec/DREIWR-2>
- Dreyfus, H. L. (2005). Overcoming the Myth of the Mental: How Philosophers Can profit from the Phenomenology of Everyday Expertise. *Proceedings and Addresses of the APA*, 79(2), 47–65. <http://link.springer.com/10.1007/s11245-006-0006-1>
- Dreyfus, H. L. (2007). The return of the myth of the mental. *Inquiry*, 50(4), 352–365.
<https://doi.org/10.1080/00201740701489245>
- Egginton, R., Newman, C., Walsh, M., Jenkinson, A., Bennett, D., & O'Brien, T. (2006). Kinematic characteristics of Achondroplasia. *Gait & Posture*, 24S, S249–S250.
<https://doi.org/10.1016/j.gaitpost.2006.11.170>

- Eich, E., Reeves, J. L., Jaeger, B., & Graff-Radford, S. B. (1985). Memory for pain: Relation between past and present pain intensity. *Pain*, 23(4), 375–380. [https://doi.org/10.1016/0304-3959\(85\)90007-7](https://doi.org/10.1016/0304-3959(85)90007-7)
- Ericsson, K. A. (2006). The influence of experience and deliberate practice on the development of superior expert performance. In the Cambridge Handbook of expertise and expert performance. In *The Cambridge Handbook of Expertise and Expert Performance* (pp. 685–706).
- Fitts, P. M., & Posner, M. I. (1967). *Human Performance*. Brooks.
- Ford, P., Hodges, N. J., & Williams, A. M. (2005). Online Attentional-Focus Manipulations in a Soccer-Dribbling Task : Implications for the Proceduralization of Motor Skills. *Journal of Moto*, 37(5), 386–394. <https://doi.org/10.3200/JMBR.37.5.386-394>
- Fox, A. T., Palmer, R. D., & Davies, P. (2002). Do “Shufflebottoms” bottom shuffle? *Archives of Disease in Childhood*, 87(6), 552–554. <https://doi.org/10.1136/ad.87.6.552>
- Freedland, R. L., & Bertenthal, B. I. (1994). Developmental Changes in Interlimb Coordination: Transition to Hands-and-Knees Crawling. *Psychological Science*, 5(1), 26–32.
- Fridland, E. (2013). *Imitation, Skill Learning, and Conceptual Thought: An Embodied, Developmental Approach*. 203–224. https://doi.org/10.1007/978-94-007-5419-5_10
- Fridland, E. (2014). They’ve lost control: Reflections on skill. *Synthese*, 191(12), 2729–2750. <https://doi.org/10.1007/s11229-014-0411-8>
- Gallagher, S. (1986). Body Image and Body Schema : A Conceptual Clarification. *The Journal of Mind and Behavior*, 7(4), 541–554.
- Gallagher, S., & Cole, J. (1995). Body Image and Body Schema in a Deafferented Subject. *The Journal of Mind and Behavior*, 16(4), 369–389.
- Gallagher, S., & Zahavi, D. (2021). *The Phenomenological Mind* (Third). Routledge.
- Geber, M., & Dean, R. F. A. (1957). Gesell Tests on African Children. *Pediatrics*, 1055–1065.

- Geeves, A., Christensen, W., Sutton, J., & McIlwain, D. (2008). Roger Chaffin, Gabriela Imreh & Mary Crawford, *Practicing Perfection: Memory and Piano Performance*. New York: Laurence Erlbaum Associates, 2002. ISBN 0-80-582610-6. *Empirical Musicology Review*, 3(3), 163–172.
- Geeves, A., McIlwain, D. J. F., Sutton, J., & Christensen, W. (2014). To Think or Not To Think: The apparent paradox of expert skill in music performance. *Educational Philosophy and Theory*, 46(6), 674–691.
<https://doi.org/10.1080/00131857.2013.779214>
- Gessell, A., & Amatruda, C. S. (1941). *Developmental diagnosis; normal and abnormal child development*. Hoeber.
- Gessell, A., Thompson, H., & Amatruda, C. S. (1934). *Infant behavior: Its genesis and growth*.
- Gettier, E. (1966). Is Justified True Belief Knowledge? *Analysis*, 23, 1–3.
- Gibson, E. (1988). Exploratory Behavior In The Development Of Perceiving, Acting, And The Acquiring Of Knowledge. *Annual Review of Psychology*, 39(1), 1–41.
<https://doi.org/10.1146/annurev.psych.39.1.1>
- Gibson, E. J., & Anne D. Pick. (2000). *An Ecological Approach to Perceptual Learning and Development*. Oxford University Press.
- Gibson, E. J., Riccio, G., Schmuckler, M. A., Stoffregen, T. A., Rosenberg, D., & Taormina, J. (1987). Detection of the Traversability of Surfaces by Crawling and Walking Infants. *Journal of Experimental Psychology: Human Perception and Performance*, 13(4), 533–544. <https://doi.org/10.1037/0096-1523.13.4.533>
- Gibson, E. J., & Walk, R. D. (1960). The “Visual Cliff”. *Scientific American*, 202(4), 64–71. <https://doi.org/10.1038/scientificamerican0460-64>
- Gill, S. V., Adolph, K. E., & Vereijken, B. (2009). Change in action: How infants learn to walk down slopes. *Developmental Science*, 12(6), 888–902.
<https://doi.org/10.1111/j.1467-7687.2009.00828.x>
- Gilmore, R. O., Baker, T. J., & Grobman, K. H. (2004). Stability in Young Infants’ Discrimination of Optic Flow. *Developmental Psychology*, 40(2), 259–270.
<https://doi.org/10.1037/0012-1649.40.2.259>
- Gilmore, R. O., & Rettke, H. J. (2003). Four-month-olds’ discrimination of optic flow patterns depicting different directions of observer motion. *Infancy*, 4(2), 177–200.
https://doi.org/10.1207/S15327078IN0402_02

- Godden, D. R., & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: on land and underwater. *British Journal of Psychology*, 66(3), 325–331.
- Goethals, G. R., & Reckman, R. F. (1973). The perception of consistency in attitudes. *Journal of Experimental Social Psychology*, 9(6), 491–501.
[https://doi.org/10.1016/0022-1031\(73\)90030-9](https://doi.org/10.1016/0022-1031(73)90030-9)
- Goldman, A. I. (1992). Empathy , Mind , and Morals. *Proceedings and Addresses of the American Philosophical Association*, 66(3), 17–41.
- Goodale, M. A., & Milner, A. D. (2010). Two visual streams: Interconnections do not imply duplication of function. *Cognitive Neuroscience*, 1(1), 65–68.
<https://doi.org/10.1080/17588920903511635>
- Gopnik, A., Meltzoff, A. N., & Kuhl, P. K. (1999). *The Scientist in the Crib: What Early Learning Tells Us About the Mind*. Harper Collins e-books.
- Gould, S. J. (2004). *Triumph and Tragedy in Mudville: A Lifelong Passion for Baseball*. W. W. Norton & Co.
- Gray, R. (2004). Attending to the Execution of a Complex Sensorimotor Skill : Expertise Differences , Choking , and Slumps. *Journal of Experimental Psychology: Applied*, 10(1), 42–54. <https://doi.org/10.1037/1076-898X.10.1.42>
- Grillner, S., & Wallen, P. (2004). Innate versus learned movements - a false dichotomy? *Progress in Brain Research*, 143, 1–12.
- Gucciardi, D. F., & Dimmock, J. A. (2008). Choking under pressure in sensorimotor skills: Conscious processing or depleted attentional resources? *Psychology of Sport and Exercise*, 9(1), 45–59. <https://doi.org/10.1016/j.psychsport.2006.10.007>
- Gustafson, G. E. (1984). Effects of the ability to locomote on infants' social and exploratory behaviors: An experimental study. *Developmental Psychology*, 20(3), 397–405. <https://doi.org/10.1037/0012-1649.20.3.397>
- Haehl, V., Vardaxis, V., & Ulrich, B. (2000). Learning to cruise: Bernstein's theory applied to skill acquisition during infancy. *Human Movement Science*, 19(5), 685–715. [https://doi.org/10.1016/S0167-9457\(00\)00034-8](https://doi.org/10.1016/S0167-9457(00)00034-8)
- Haugeland, J. (1998). *Having Thought: Essays in the Metaphysics of Mind*. Harvard University Press.

- Heglund, N. C., Willems, P. A., Penta, M., & Cavagna, G. A. (1995). Energy-saving gait mechanics with head-supported loads. *Nature*, 375(6526), 52–54. <https://doi.org/10.1038/375052a0>
- Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56(5), 872–876. <https://doi.org/10.1037/h0040546>
- Henson, R. N., & Gagnepain, P. (2010). Predictive, interactive multiple memory systems. *Hippocampus*, 20, 1315–1326. <https://doi.org/10.1002/hipo.20857>
- Herbert Leiderman, P., Babu, B., Kagia, J., Kraemer, H. C., & Leiderman, G. F. (1973). African infant precocity and some social influences during the first year. *Nature*, 242(5395), 247–249. <https://doi.org/10.1038/242247a0>
- Hoerl, C. (2017). Memory and the Concept of Time. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory* (pp. 207–218). Routledge.
- Holler, J., Turner, K., & Varcianna, T. (2013). It ' s on the tip of my fingers : Co-speech gestures during lexical retrieval in different social contexts. *LANGUAGE AND COGNITIVE PROCESSES*, 28(10), 1509–1518. <https://doi.org/10.1080/01690965.2012.698289>
- Hopkins, B., & Westra, T. (1989). Maternal Expectations of Their Infants' Development: Some Cultural Differences. *Developmental Medicine & Child Neurology*, 31(3), 384–390. <https://doi.org/10.1111/j.1469-8749.1989.tb04008.x>
- Hopkins, B., & Westra, T. (1990). Motor development, maternal expectations, and the role of handling. *Infant Behavior and Development*, 13(1), 117–122. [https://doi.org/10.1016/0163-6383\(90\)90011-V](https://doi.org/10.1016/0163-6383(90)90011-V)
- Husserl, E. (1991). *On the phenomenology of the consciousness of internal time (1893-1917)* (J. B. Brough (Ed.)). Kluwer Academic Publishers.
- Hutto, D. D. (2017). Memory and Narrativity. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory* (pp. 192–204). Routledge.
- Hutto, D. D., & Myin, E. (2013). *Radicalizing Enactivism: Basic Minds without Content*. MIT Press.
- Hutto, D. D., & Myin, E. (2017). *Evolving Enactivism: Basic Minds Meet Content*. MIT Press.

- Hutto, D. D., & Peeters, A. (2018). The Roots of Remembering Radically Enactive Recollecting. In K. Michaelian (Ed.), *New Directions in the Philosophy of Memory* (pp. 97–118). Routledge.
- Iloeje, S. O. L., Obiekwe, V. U., & Kaine, W. N. (1991). Gross motor development of Nigerian children. *Annals of Tropical Paediatrics*, 11(1), 33–39.
<https://doi.org/10.1080/02724936.1991.11747475>
- Ivanhoe, P. J., & Norden, B. W. Van (Eds.). (2005). *Readings in Classical Chinese Philosophy* (Second). Hackett.
- Jenkins, S. (2007). The Use of Swing Keys by Elite Tournament Professional Golfers. *International Journal of Sports Science & Coaching*, 2(1_suppl), 199–229.
<https://doi.org/10.1260/174795407789705460>
- Joh, A. S., & Adolph, K. E. (2006). Learning from falling. *Child Development*, 77(1), 89–102. <https://doi.org/10.1111/j.1467-8624.2006.00858.x>
- Karasik, L. B., Tamis-Lemonda, C. S., & Adolph, K. E. (2011). Transition from crawling to walking and infants' actions with objects and people. *Child Development*, 82(4), 1199–1209. <https://doi.org/10.1111/j.1467-8624.2011.01595.x>
- Karasik, L. B., Tamis-LeMonda, C. S., Adolph, K. E., & Dimitropoulou, K. A. (2008). How mothers encourage and discourage infants' motor actions. *Infancy*, 13(4), 366–392. <https://doi.org/10.1080/15250000802188776>
- Kelly, M. R. (2015). The Subject as Time: Merleau-Ponty's transition from Phenomenology to Ontology. In D. Morris & K. Maclaren (Eds.), *Time, Memory, Institution: Merleau-Ponty's New Ontology of Self* (pp. 199–215). Ohio University Press.
- Kelly, S. D., & Barr, D. J. (1999). Offering a Hand to Pragmatic Understanding : The Role of Speech and Gesture in Comprehension and Memory. *Journal of Memory and Language*, 40, 577–592.
- Kermoian, R., & Campos, J. J. (1988). Locomotor Experience : A Facilitator of Spatial Cognitive Development Author (s): Rosanne Kermoian and Joseph J . Campos
 Published by : Wiley on behalf of the Society for Research in Child Development
 Stable URL : <http://www.jstor.org/stable/1130258> REF. *Child Development*, 59(4), 908–917.
- Kilbride, J. E., Robbins, M. C., & Kilbride, P. L. (1970). The Comparative Motor Development of Baganda, American White, and American Black Infants. *American Anthropologist*, 72, 1422–1428.

- Klopfer, P. (1988). Metaphors for development: How important are experiences early in life? *Developmental Psychobiology*, 21(7), 671–678.
<https://doi.org/10.1002/dev.420210706>
- Kuklina, E. V., Ramakrishnan, U., Stein, A. D., Barnhart, H. H., & Martorell, R. (2004). Growth and diet quality are associated with the attainment of walking in rural Guatemalan infants. *Journal of Nutrition*, 134(12), 3296–3300.
<https://doi.org/10.1093/jn/134.12.3296>
- Lakoff, G., & Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press.
- Ledebt, A., & Bril, B. (2000). Acquisition of upper body stability during walking in toddlers. *Developmental Psychobiology*, 36(4), 311–324.
[https://doi.org/10.1002/\(SICI\)1098-2302\(200005\)36:4<311::AID-DEV6>3.0.CO;2-V](https://doi.org/10.1002/(SICI)1098-2302(200005)36:4<311::AID-DEV6>3.0.CO;2-V)
- Levine, L. J. (1997). Reconstructing Memory for Emotions. *Journal of Experimental Psychology: General*, 126(2), 165–177. <https://doi.org/10.1037/0096-3445.126.2.165>
- Loader, P. (2013). Is my Memory an Extended Notebook ? *Review of Philosophy and Psychology*, 4, 167–184. <https://doi.org/10.1007/s13164-012-0123-2>
- Locke, J. (n.d.). *An Essay Concerning Human Understanding*.
- Lockhart, R. S., Craik, F. I. M., & Jacoby, L. (1976). Depth of processing, recognition, and recall. In J. Brown (Ed.), *Recall and recognition*. Wiley.
- Loftus, E. F. (1974). Reconstructing memory: The incredible eyewitness. *Jurimetrics Journal*, 15(3), 188–193.
- Loftus, E. F., & Pickrell, J. E. (1995). The Formation of False Memories. In *Psychiatric Annals* (Vol. 25, Issue 12, pp. 720–725). <https://doi.org/10.3928/0048-5713-19951201-07>
- Luria, A. R. (1973). *The Working Brain: An Introduction to Neuropsychology* (B. Haigh (Ed.)). Basic Books.
- MacPherson, A., Collins, D., & Morriss, C. (2008). Is what you think what you get? Optimizing mental focus for technical performance. *Sport Psychologist*, 22(3), 288–303. <https://doi.org/10.1123/tsp.22.3.288>
- Mahler, M. S., Pine, F., & Bergman, A. (1975). *The Psychological Birth of the Human Infant: Symbiosis and Individuation*. Karnac.

- Maloiy, G. M. O., Heglund, N. C., Prager, L. M., Cavagna, G. A., & Taylor, C. R. (1986). Energetic cost of carrying loads: have African women discovered an economic way? *Nature*, 319, 668–669.
- Martin, C. B., & Deutscher, M. (1966). Remembering. *The Philosophical Review*, 75(2), 161–196.
- McGraw, M. B. (1945). *The neuromuscular maturation of the human infant*. Columbia University Press.
- Merleau-Ponty, M. (2012). *Phenomenology of Perception*. Routledge.
- Metcalf, J. S., & Clark, J. E. (2000). Sensory information affords exploration of posture in newly walking infants and toddlers. *Infant Behavior and Development*, 23(3–4), 391–405. [https://doi.org/10.1016/S0163-6383\(01\)00054-6](https://doi.org/10.1016/S0163-6383(01)00054-6)
- Michaelian, K. (2011). Generative memory. *Philosophical Psychology*, 24(3), 323–342. <https://doi.org/10.1080/09515089.2011.559623>
- Michaelian, K. (2012). Is external memory memory ? Biological memory and extended mind. *Consciousness and Cognition*, 21, 1154–1165. <https://doi.org/10.1016/j.concog.2012.04.008>
- Michaelian, K. (2016a). Against Discontinuism. *Seeing the Future*, 62–92. <https://doi.org/10.1093/acprof:oso/9780190241537.003.0004>
- Michaelian, K. (2016b). Confabulating, misremembering, relearning: The simulation theory of memory and unsuccessful remembering. *Frontiers in Psychology*, 7(NOV), 1–13. <https://doi.org/10.3389/fpsyg.2016.01857>
- Michaelian, K. (2016c). Mental time travel: Episodic memory and our knowledge of the personal past. In *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. MIT Press. <https://doi.org/10.1215/00318108-3878563>
- Michaelian, K., & Sutton, J. (2017). Memory. In *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2017/entries/memory/>
- Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford University Press.
- Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, 46(3), 774–785. <https://doi.org/10.1016/j.neuropsychologia.2007.10.005>

- Milner, B. (1962). Les troubles de la memoire accompagnant des lesions hippocampiques bilaterales. In *Physiologie de l'hippocampe* (pp. 257–272). C.N.R.S.
- Miracchi, L. (2015). Competence to know. *Philosophical Studies*, 172, 29–56.
<https://doi.org/10.1007/s11098-014-0325-9>
- Moessinger, A. C. (1983). Fetal akinesia deformation sequence: An animal model. *Pediatrics*, 72(6), 857–863. <https://doi.org/10.1542/peds.72.6.857>
- Montero, B. (2010). Does bodily awareness interfere with highly skilled movement? *Inquiry*, 53(2), 105–122. <https://doi.org/10.1080/00201741003612138>
- Morris, D., & Maclaren, K. (2015). Introduction. In D. Morris & K. Maclaren (Eds.), *Time, Memory, Institution: Merleau-Ponty's New Ontology of Self* (pp. 1–25). Ohio University Press.
- Morsella, E., & Krauss, R. M. (2002). *Movement Facilitates Speech Production: A Gestural Feedback Model*. 212, 1–33.
<http://www.columbia.edu/~rmk7/PDF/Mvmt.pdf>
- Morsella, E., & Krauss, R. M. (2004). The role of gestures in spatial working memory and speech. *American Journal of Psychology*, 117(3), 411–424.
<https://doi.org/10.2307/4149008>
- Muldoon, M. S. (2006). *Tricks of Time: Bergson, Merleau-Ponty and Ricoeur in Search of Time, Self and Meaning*. Duquesne University Press.
- Mullally, S. L., & Maguire, E. A. (2014). Memory, imagination, and predicting the future: A common brain mechanism? *Neuroscientist*, 20(3), 220–234.
<https://doi.org/10.1177/1073858413495091>
- Mullen, R., & Hardy, L. (2010). Conscious processing and the process goal paradox. *Journal of Sport and Exercise Psychology*, 32(3), 275–297.
<https://doi.org/10.1123/jsep.32.3.275>
- Myin, E., & Zahidi, K. (2015). The Extent of Memory: From Extended to Extensive Mind. In A. Coliva, V. Munz, & D. Moyal-Sharrock (Eds.), *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium* (pp. 391–408). De Gruyter.
- Nakayama, K., Nishimaru, H., & Kudo, N. (2001). Developmental changes in 5-hydroxytryptamine-induced rhythmic activity in the spinal cord of rat fetuses in vitro. *Neuroscience Letters*, 307(1), 1–4. [https://doi.org/10.1016/S0304-3940\(01\)01913-9](https://doi.org/10.1016/S0304-3940(01)01913-9)

- Neisser, U. (1967). *Cognitive psychology*. Appleton-Century-Crofts.
- Neisser, Ulric. (1996). Remembering as Doing. *Behavioral and Brain Sciences*, 19(2), 203–204.
- Newcombe, N., & Huttenlocher, J. (2000). *Making space: The development of spatial representation and reasoning*. MIT Press.
- Noe, A. (2004). *Action in Perception*. MIT Press.
- Noe, A. (2005). Against Intellectualism. *Analysis*, 65(4), 278–290.
- Norman, K. A., & Schacter, D. L. (1997). False recognition in younger and older adults : Exploring the characteristics of illusory memories. *Memory and Cognition*, 25(6), 838–848.
- O'Brien, G. J. (1991). Is Connectionism Commonsense? *Philosophical Psychology*, 4(2), 165–178.
- Papineau, D. (2013). In the zone. *Royal Institute of Philosophy Supplement*, 73, 175–196. <https://doi.org/10.1044/leader.GL.22122017.12>
- Papineau, D. (2015). Choking and The Yips. *Phenomenology and the Cognitive Sciences*, 14, 295–308. <https://doi.org/10.1007/s11097-014-9383-x>
- Pavese, C. (2016). Skill in epistemology I : Skill and knowledge. *Philosophy Compass*, 11, 642–649. <https://doi.org/10.1111/phc3.12359>
- Perri, T. (2017). Henri Bergson. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory*2 (pp. 510–518). Routledge.
- Perrin, D. (2011). Une défense de l'approche simulationniste du souvenir épisodique. *Dialogue-Canadian Philosophical Review*, 50(1), 39–76. <https://doi.org/10.1017/S0012217311000114>
- Perrin, D. (2018). A Case for Procedural Causality in Episodic Recollection. In K. Michaelian, D. Debus, & D. Perrin (Eds.), *New Directions in the Philosophy of Memory* (1st ed., pp. 33–51). Routledge.
- Piaget, J. (1952). *The Origins of Intelligence in Children*. International Universities Press.
- Piontelli, A. (2010). Development of normal fetal movements: The first 25 weeks of gestation. In *Development of Normal Fetal Movements: The First 25 Weeks of Gestation*. <https://doi.org/10.1007/978-88-470-1402-2>

Plato. (n.d.). *Theaetetus*.

Poidevin, R. Le. (2017). Memory and the Metaphysics of Time. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory* (pp. 219–227). Routledge.

Poldrack, R. A., Sabb, F. W., Foerde, K., Tom, S. M., Asarnow, R. F., Bookheimer, S. Y., & Knowlton, B. J. (2005). The neural correlates of motor skill automaticity. *Journal of Neuroscience*, 25(22), 5356–5364.
<https://doi.org/10.1523/JNEUROSCI.3880-04.2005>

Pollard, B. (2006). Explaining Actions with Habits. *American Philosophical Quarterly*, 43(1), 57–69.

Proust, M. (2003). *The Way by Swann's* (C. Prendergast (Ed.)). Penguin.

Pyers, J. E., Magid, R., Gollan, T. H., & Emmorey, K. (2021). Gesture Helps , Only If You Need It : Inhibiting Gesture Reduces Tip-of-the-Tongue Resolution for Those With Weak Short-Term Memory. *Cognitive Science*, 45.
<https://doi.org/10.1111/cogs.12914>

Pylyshyn, Z. W. (2003). *Seeing and Visualizing: It's Not What You Think*. MIT Press.

Rabain-Jamin, J., & Wornham, W. L. (1993). Among West Africans in Paris ? *Early Development and Parenting*, 2(February 1992), 107–119.

Ren, Y., & Argote, L. (2011). Transactive Memory Systems 1985 – 2010 : An Integrative Framework of Key Dimensions , Antecedents , and Consequences. *The Academy of Management Annals*, 5(1), 189–229. <https://doi.org/10.1080/19416520.2011.590300>

Robins, S. (2016a). Misremembering. *Philosophical Psychology*, 29(3), 432–447.
<https://doi.org/10.1080/09515089.2015.1113245>

Robins, S. (2016b). Representing the past: memory traces and the causal theory of memory. *Philosophical Studies*, 173(11), 2993–3013.
<https://doi.org/10.1007/s11098-016-0647-x>

Robins, S. K. (2017). Memory Traces. In S. Bernecker & K. Michaelian (Eds.), *The Routledge Handbook of Philosophy of Memory*. Routledge.

Robson, P. (1984). Prewalking locomotor movements and their use in predicting standing and walking. *Child: Care, Health and Development*, 10(5), 317–330.
<https://doi.org/10.1111/j.1365-2214.1984.tb00189.x>

- Roediger, H. L. (1980). Memory metaphors in cognitive psychology. *Memory & Cognition*, 8(3), 231–246. <https://doi.org/10.3758/BF03197611>
- Roy, S., & Park, N. W. (2010). Dissociating the memory systems mediating complex tool knowledge and skills. *Neuropsychologia*, 48(10), 3026–3036. <https://doi.org/10.1016/j.neuropsychologia.2010.06.012>
- Roy, S., & Park, N. W. (2018). Response: Commentary: Effects of dividing attention on memory for declarative and procedural aspects of tools use. *Frontiers in Psychology*, 9(631), 1–3. <https://doi.org/10.3389/fpsyg.2018.00631>
- Russell, B. (2009). *Human Knowledge: Its Scope and Limits*. Routledge.
- Rust, J. (2009). John Searle. In *Angewandte Chemie International Edition*, 6(11), 951–952. (Vol. 3, Issue 1). Continuum.
- Ryle, G. (1949). *The Concept of Mind*. University of Chicago Press.
- Schacter, D. L. (2019). Implicit Memory, Constructive Memory, and Imagining the Future: A Career Perspective. *Perspectives on Psychological Science*, 14(2), 256–272. <https://doi.org/10.1177/1745691618803640>
- Schacter, D. L., & Addis, D. R. (2007a). Constructive memory: The ghosts of past and future. *Nature*, 445(7123), 27. <https://doi.org/10.1038/445027a>
- Schacter, D. L., & Addis, D. R. (2007b). The cognitive neuroscience of constructive memory : remembering the past and imagining the future. *Philosophical Transactions of the Royal Society*, 362, 773–786. <https://doi.org/10.1098/rstb.2007.2087>
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, 8(9), 657–661. <https://doi.org/10.1038/nrn2213>
- Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology*, 49, 289–318. <https://doi.org/10.1146/annurev.psych.49.1.289>
- Schacter, D. L., Reiman, E., Curran, T., Yun, L. S., Bandy, D., Mcdermott, K. B., & Roediger, H. L. (1996). Neuroanatomical Correlates of Veridical and Illusory Recognition Memory : Evidence from Positron Emission Tomography. *Neuron*, 17, 267–274.
- Schacter, D. L., & Tulving, E. (1994). *Memory Systems 1994*. MIT Press.

- Schacter, D. L., Verfaellie, M., & Pradere, D. (1996). The Neuropsychology of Memory Illusions : False Recall and Recognition in Amnesic Patients. *Journal of Memory and Language*, 35, 319–334.
- Schmuckler, M. A. (1996). Development of Visually Guided Locomotion: Barrier Crossing by Toddlers. *Ecological Psychology*, 8(3), 209–236.
https://doi.org/10.1207/s15326969eco0803_2
- Shanon, B. (1998). Metaphorical pluralism - not on the substantive level! *Behavioral and Brain Sciences*, 21(1), 164–168.
- Shanton, K., & Goldman, A. (2010). Simulation theory. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(4), 527–538. <https://doi.org/10.1002/wcs.33>
- Sheets-Johnstone, M. (2003). Kinesthetic Memory. *Theoria Et Historia Scientiarum*, VII(1), 69–92.
- Shirley, M. M. (1933). *The First Two Years: A study of twenty-five babies: Vol. I.* The University of Minnesota Press.
- Shumway-Cook, A., & Woollacott, M. H. (2012). *Motor Control: Translating Research into Clinical Practice* (4th ed.). Wolters Kluwer Health.
- Sosa, E. (2007). *A Virtue Epistemology: Apt and Reflective Knowledge* (Vol. 1). Clarendon Press.
- Sosa, E. (2010). How Competence Matters in Epistemology. *Philosophical Perspectives*, 24(Epistemology (2010)), 465–475.
- Sosa, E. (2015). The Unity of Action, Perception and Knowledge. In *Judgement and Agency* (pp. 7–34). Oxford University Press.
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). *Google Effects on Memory : Information at Our Fingertips. August*, 776–779.
- Spelke, E. S., & Newport, E. L. (1998). Nativism, empiricism, and the development of knowledge. In W. Damon & R. M. Lerner (Eds.), *Handbook of child psychology: Theoretical models of human development* (pp. 275–340). John Wiley & Sons Inc.
- Stanley, J. (2011a). *Know How*. Oxford University Press.
- Stanley, J. (2011b). Knowing (How). *Nous*, 45(2), 207–238.

- Stanley, J., & Krakauer, J. W. (2013). Motor skill depends on knowledge of facts. *Frontiers in Human Neuroscience*, 7(AUG), 1–11. <https://doi.org/10.3389/fnhum.2013.00503>
- Stanley, J., & Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98(8), 411–444.
- Stanley, J., & Williamson, T. (2017). Skill. *Nous*, 51(4), 713–726. <https://doi.org/10.1111/nous.12144>
- Stern, D. G. (1991). Models of memory: Wittgenstein and cognitive science. *Philosophical Psychology*, 4(2), 203–218. <https://doi.org/10.1080/09515089108573027>
- Suddendorf, T., & Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *The Behavioral and Brain Sciences*, 30(3), 299–313; discussion 313–351. <https://doi.org/10.1017/S0140525X07001975>
- Super, C. M. (1976). Environmental Effects on Motor Development: the Case of ‘African Infant Precocity.’ *Developmental Medicine & Child Neurology*, 18(5), 561–567. <https://doi.org/10.1111/j.1469-8749.1976.tb04202.x>
- Sutherland, D. H., Olshen, R. A., Biden, E. N., & Wyatt, M. P. (1988). *The Development of Mature Walking*. MacKeith Press.
- Sutherland, D., Olshen, R., Cooper, L., & Woo, S. (1980). The development of mature gait. *The Journal of Bone and Joint Surgery*, 62, 336–353. [https://doi.org/10.1016/S0966-6362\(97\)00029-5](https://doi.org/10.1016/S0966-6362(97)00029-5)
- Sutton, J. (n.d.). *Constructive Memory and Distributed Cognition: Towards an Interdisciplinary Framework* (pp. 290–303).
- Sutton, J. (2007). Batting, habit and memory: The embodied mind and the nature of skill. *Sport in Society*, 10(5), 763–786. <https://doi.org/10.1080/17430430701442462>
- Sutton, J. (2009). The Feel of the World: Exograms, Habits, and the Confusion of Types of Memroy. In A. Kania (Ed.), *Memento: philosophers on film* (pp. 65–86). Routledge.
- Sutton, J. (2010). Exograms and Interdisciplinarity : History , the Extended Mind , and the Civilizing Process. In R. Menary (Ed.), *The extended mind* (pp. 189–225). MIT Press.

- Sutton, J., McIlwain, D., Christensen, W., & Geeves, A. (2011). Applying Intelligence to the Reflexes: Embodied Skills and Habits Between Dreyfus and Descartes. *Journal of the British Society for Phenomenology*, 42(1), 78–103.
- Sutton, J., & Williamson, K. (2014). Embodied Remembering. In L. Shapiro (Ed.), *The Routledge Handbook of Embodied Cognition*.
- Suzuki, S., & Yamamuro, T. (1985). Fetal movement and fetal presentation. *Early Human Development*, 11(3–4), 255–263. [https://doi.org/10.1016/0378-3782\(85\)90079-9](https://doi.org/10.1016/0378-3782(85)90079-9)
- Szpunar, K. K., Spreng, R. N., & Schacter, D. L. (2014). A taxonomy of prospection : Introducing an organizational framework for future- oriented cognition. *Proceedings of the National Academy of Sciences of the United States of America*, 111(52), 18414–18421. <https://doi.org/10.1073/pnas.1417144111>
- Tamis-LeMonda, C. S., Adolph, K. E., Dimitropoulou, K. A., & Zack, E. (2007). “No! Don’t! Stop!”: Mothers’ Words for Impending Danger. *Parenting*, 7(1), 1–25. <https://doi.org/10.1080/15295190709336774>
- Taylor, C. R. (1995). Freeloading women. *Nature*, 375(6526), 17. <https://doi.org/10.1038/375017a0>
- Taylor, C. R., Heglund, N. C., McMahon, T. A., & Looney, T. R. (1980). Energetic Cost of Generating Muscular Force During Running: A Comparison of Large and Small Animals. *Journal of Experimental Biology*, 86(1), 9–18. <https://doi.org/10.1242/jeb.86.1.9>
- Thelen, E., Fisher, D. M., & Ridley-Johnson, R. (1984). The relationship between physical growth and a newborn reflex. *Infant Behavior and Development*, 7, 479–493.
- Thelen, Esther. (1979). Rhythmical stereotypies in normal human infants. *Animal Behaviour*, 27(PART 3), 699–715. [https://doi.org/10.1016/0003-3472\(79\)90006-X](https://doi.org/10.1016/0003-3472(79)90006-X)
- Thelen, Esther. (1995). Motor development: A new synthesis. *American Psychologist*, 50(2), 79–95. <https://doi.org/10.1037/0003-066X.50.2.79>
- Thelen, Esther. (2000). Grounded in the world: Developmental origins of the embodied mind. *Infancy*, 1(1), 3–28. <https://doi.org/10.4324/9780203809778>
- Thelen, Esther. (2005). Dynamic systems theory and the complexity of change. *Psychoanalytic Dialogues*, 15(2), 255–283. <https://doi.org/10.1080/10481881509348831>

- Thelen, Esther, & Fisher, D. M. (1982). Newborn Stepping: An Explanation for a “Disappearing” Reflex. *Developmental Psychology*, 18(5), 760–775.
- Thelen, Esther, & Smith, L. B. (1996). *A Dynamic Systems Approach to the Development of Cognition and Action* (Vol. 13, Issue April). MIT Press.
- Toner, J., Montero, B. G., & Moran, A. (2015). Considering the role of cognitive control in expert performance. *Phenomenology and the Cognitive Sciences*, 14, 1127–1144. <https://doi.org/10.1007/s11097-014-9407-6>
- Toner, J., Montero, B. G., & Moran, A. (2016). Reflective and prereflective bodily awareness in skilled action. *Psychology of Consciousness: Theory, Research, and Practice*, 3(4), 303–315. <https://doi.org/10.1037/cns0000090>
- Treppien, A. W. (1900). Creeping and Walking. *The American Journal of Psychology*, 12(1), 1–57.
- Tulving, E. (1983). *Elements of Episodic Memory*. Oxford University Press.
- Tulving, E., & Madigan, S. A. (1970). Memory and Verbal Learning. In *Annual Review of Psychology* (pp. 437–484).
- Ulrich, B. D. (1989). Development of stepping patterns in human infants: A dynamical systems perspective. *Journal of Motor Behavior*, 21(4), 392–408. <https://doi.org/10.1080/00222895.1989.10735491>
- Vallacher, R. R., & Wegner, D. M. (1987). Action Identification and Human Behavior. *Psychological Review*, 94(1), 3–15.
- Vallier, R. (2015). Memory - of the Future: Institution and Memory in the later Merleau-Ponty. In D. Morris & K. Maclaren (Eds.), *Time, Memory, Institution: Merleau-Ponty's New Ontology of Self* (pp. 109–129). Ohio University Press.
- Van der Meer, A. L. H. (1997). Visual guidance of passing under a barrier. *Infant and Child Development*, 6(3–4), 149–157. [https://doi.org/10.1002/\(sici\)1099-0917\(199709/12\)6:3/4<149::aid-edp154>3.0.co;2-2](https://doi.org/10.1002/(sici)1099-0917(199709/12)6:3/4<149::aid-edp154>3.0.co;2-2)
- Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the Control Homunculi in Studies of Action Control and Behavior Change. *Perspectives on Psychological Science*, 9(5), 497–524. <https://doi.org/10.1177/1745691614526414>
- Vereijken, B., Bril, B., & Ledebt, A. (1998). Development of independent walking: The critical role of double support. In *Advances in perception-action coupling* (pp. 179–183). Editions EDK.

- Vereijken, B., & Waardenburg, M. (1996). Changing patterns of interlimb coordination from supported to independent walking. *Infant Behavior and Development*, 19.
- Vereijken, Beatrix, & Adolph, K. (1999). *Transitions in the development of locomotion* (pp. 137–149). <https://doi.org/10.4219/gct-2005-168>
- Vereijken, Beatrix, & Albers, R. (1998). The nature of support in supported walking. *Infant Behavior and Development*, 21.
- Vygotsky, L. S. (1978). *Mind in Society: The Development of Higher Psychological Processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman (Eds.)). Harvard University Press.
- Waddington, C. H. (1957). *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. Routledge.
- Wechsler, D. (1963). Engrams, memory storage, and mnemonic coding. *American Psychologist*, 18(3), 149–153. <https://doi.org/10.1037/h0046456>
- Wegner, D. M., Erber, R., & Raymond, P. (1991). *Transactive Memory in Close Relationships*. 6, 923–929.
- Werner, E. E. (1972). Infants around the world: cross-cultural studies of psychomotor development from birth two years. *Journal of Cross-Cultural Psychology*, 3(2), 111–134.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, T. (2007). *The philosophy of philosophy*. Blackwell. <https://doi.org/10.4324/9780203029473>
- Woollacott, M. H., Shumway-Cook, A., & Williams, H. G. (1989). The development of posture and balance control in children. In *Development of posture and gait across the life span* (pp. 77–96).
- Wu, W. (2014). *Attention* (New Proble). Routledge.
- Xie, Q., & Young, M. E. (1999). *Integrated Child Development in Rural China* (Issue January).