PERCEPTION AND SKILL: THEORETICAL FOUNDATIONS FOR A SCIENCE OF PERCEPTION

by

ELLEN FRIDLAND

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2010

This manuscript has been read and accepted for the
Graduate Faculty in Philosophy in satisfaction of the
dissertation requirement for the degree of Doctor of Philosophy.

Date	Chair of Examining Committee Nickolas Pappas
Date	Executive Officer Iakovos Vasiliou

Michael Devitt

Alva Noë

Nickolas Pappas

Jesse Prinz

David M. Rosenthal (advisor)

Supervision Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

PERCEPTION AND SKILL: THEORETICAL FOUNDATIONS FOR A SCIENCE OF PERCEPTION

by

Ellen Fridland

Advisor: Professor David M. Rosenthal

I argue in my dissertation that if diachronic cognitive penetration is caused by skill, then such changes in perceptual processing are legitimate instances of cognitive penetration. As such, perceptual processing is not modular. I argue this by (1) presenting a detailed analysis of the definition of cognitive penetration, (2) arguing that propositional knowledge cannot account for practical know-how, and (3) providing a definition of skill that highlights its practical and irreducibly cognitive nature. Taken together, these considerations amount to an argument for the possibility of a genuine instance of cognitive penetration, which results from the regular instantiation of skill.

I could not have written this dissertation without the love, support, and encouragement of many people. I would like to thank my parents, Anatoly and Ida Fridland, for instilling in me a respect for education, a love of great books, and a fascination with extraordinary ideas; my husband, Andrew Porter, for his unwavering faith in my abilities and for his time spent editing "in the trenches." I have taken this journey with many great friends and colleagues. In particular, I'd like to thank Bana Bashour, Damien Dupont, Kristian Kemtrup, Michael Benitez, Maria Brincker, David Pereplyotchik, James Dow, Chris Sula, Arindam Chakrabarti, Amy Donahue, Susana Montesinos, Kelly Cutler, Orrie King, Sebastian Frank, Ruthie Ackerman, and Gene and Jen Fridland for countless hours of discussion, humor, and kindness.

TABLE OF CONTENTS

CHAPTER 1-
Perception and Cognition: Complications and Motivations

1. Introduction

There are some human activities that clearly depend on cognition for their successful instantiation. Reading and arithmetic are two such examples. There are other activities, like playing the piano and doing gymnastics, which may be less obviously dependent on cognition for their existence, but are clearly connected to experience and skill. Still, there are other activities, such as digestion and circulation, where belief and experience seem to be utterly irrelevant.

Perception, unfortunately, does not neatly fall into any of the above categories. It is not like thinking, because it does not involve beliefs or desires and it is not like skill, since one need not practice or train in order to learn how to do it. But it isn't like digestion, either, since perception has an obvious connection to beliefs and behavior. Trivially, at least, many of our beliefs depend on the inputs of our perceptual experiences. But is the relationship between cognition and perception symmetrical? That is, does perception depend on beliefs and knowledge like knowledge and beliefs depend on perception? Can my thinking about red change the way my perceptual system processes red sensations? This is the question of cognitive penetration and this is the question that I explore in this dissertation.

Now, most everyone agrees that normal, mature, healthy, human adult perception is intimately connected to cognition. That is, mature, healthy human adults not only see things, but they see things *as* things. For instance, when you hear a horn blaring, you not only hear some raw sound, but you also have thoughts, assumptions, and expectations accompanying that sound. You may think, either implicitly or explicitly: "that is a horn blaring," or "horns sound different

than bells," or "horns come from cars," or "that blaring horn is coming from the street outside my window," or "there must be traffic on the street," and so on. You not only hear the horn, but you place the horn into the logical space of reasons.

The main sticking point in debates about the cognitive penetration of perception is not whether perception is usually *accompanied by* cognition, but rather, whether conceptual, cognitive and rational structures penetrate all the way down, *impacting, influencing, structuring or constituting* the sensory core of a perceptual event. That cognition affects the intentional aspects of a perceptual state is not controversial; what is controversial, however, is the claim that cognition impacts the qualitative character of a perceptual event.

There are many issues related to the cognitive penetration of perception that I will do my best to keep distinct. In this dissertation, I do not address how perception informs cognition or belief formation or how the relationship between cognition and perception affects the possibility of knowledge. I am not concerned with questions of whether perception has conceptual or nonconceptual content or if this content is internal or external. I do not engage in the debate between the representationalists and the qualia realists. Nor do I discuss whether it is best to think of perceptual properties as being processed by computational functions or by connectionist networks.

This is my task: I attempt to ascertain if the diachronic cognitive penetration of perception by skill would count as a legitimate instance of cognitive penetration. I do this in four main chapters: in chapter two, I examine what it would take for any event or state to cognitively penetrate perception; in chapter three, I defend the idea that skill is not reducible to propositional knowledge; in chapters four and five, I argue that skill is a legitimate instance of cognition, such

that the proper relationship to it, by perception, would constitute an instance of cognitive penetration.

In many ways, this dissertation can be seen as an attempt to elucidate the conditions that must obtain if the claims of enactive perception are to entail the cognitive penetration of perception. Enactive perception theorists such as Alva Noë and Susan Hurley claim that sensorimotor skill is constitutive of the qualitative character of perceptual events. On their view, understanding sensorimotor contingencies is constitutive of the way the world appears.[1] That is, what we know about these contingencies determines what we perceive. Enactive perception theorists take themselves to be committed to a strong form of cognitive penetration.

In this dissertation, I do not argue in support of the theory of enactive perception, but rather, I lay out the theoretical conditions which must obtain for skill *of any kind* to cause the cognitive penetration of perceptual processing. These conditions will quite clearly apply to sensorimotor skill. Enactive perception, however, is only my starting point—it will not be my central concern.

Before beginning my argument, I'd like to motivate the problem of cognitive penetration with some general considerations. Pre-theoretically, it seems perfectly natural to think of sensing, or the production of qualitative character, as something that happens in a purely mechanical fashion. Like digestion, it would seem that certain stimuli trigger fixed physiological responses and produce corresponding hardwired results. If sensing is like digestion, then one

[1] "The world makes itself available to the perceiver through physical movement and interaction. In this book I argue that all perception is touch-like in this way: Perceptual experience acquires content thanks to our possession of bodily skills. *What we perceive* is determined by *what we do* (or what we know how to do); it is determined by what we are *ready* to do. In ways I try to make precise, we *enact* our perceptual experience; we act it out. To be a perceiver is to understand, implicitly, the effects of movement on sensory stimulation…*the enactive approach* is that our ability to perceive not only depends on, but is constituted by, our possession of this sort of sensorimotor knowledge" (italics in original) Alva Noë, *Action in Perception* (Cambridge, MA: MIT Press, 2004), 1. See also Susan Hurley, *Consciousness in Action* (Cambridge, MA: Harvard University Press, 1998).

may wonder why, if thinking about our breakfast does not affect the way we digest our bagels, should thinking about the world influence our sensation of blue or loud or hot?

As empiricists, it is perfectly natural to think of sensations as qualities about which we reason, but not themselves the products of our reasoning. This is, after all, the position of our empiricist forefathers, John Locke, George Berkeley and David Hume.[2] They took it that the senses provide us with the raw materials out of which to fashion ideas and experiences. Sensations are not impacted by thoughts, though our experiences of them may be by way of the judgments we make about them. That is, experience is not of raw sensations, but raw sensations are the basic inputs about which judgments are formed and out of which experiences are constructed.

This makes sense, for if the qualitative character of a perceptual event was determined by what we believed, then things should look and sound and taste much different than they do. This must be the case since I know all sorts of things about perceptual illusions but that does not make those illusions disappear. I know that a pencil in water is not bent, but my knowledge does not change the fact that it looks bent; I know that thunder occurs at the same time as lightning, but I still hear it only after seeing the lightning with which it coincides; and I know the Earth is not flat, but the horizon always looks as though, if I just travelled far enough, eventually I'd fall off. If perception were cognitively penetrable, then these beliefs should affect my sensory experiences of them. On the contrary, however, it seems that things appear the same no matter what I believe.

---

[2] See John Locke, *An Essay on Human Understanding* (London: Penguin Books, 1997); George Berkeley, *A Treatise Concerning the Principles of Human Knowledge* (USA: Oxford University Press, 1998); David Hume, *An Enquiry Concerning Human Understanding* (Indianapolis, IN: Bobbs- Merrill, 1955).

So, why should the cognitive penetration of perception be an issue at all? The problems begin when we begin to think about expertise, perceptual learning, Gestalt shifts, the impacts of mood on judging perceptual qualities, and the anatomy of the brain.

Perceptual learning should make us pause before declaring that perception is immune to the impacts of cognition because if one can learn to discriminate new or different qualities as a result of training, then it would seem that beliefs, knowledge, and experience do affect perception. The wine taster is able to discriminate perceptual nuances that the layperson cannot. The same goes for the radiologist. It seems that artists are more sensitive to color, musicians to sound, gourmands to taste, and athletes to the positions of their bodies. Expertise, it seems, is intimately related to the sensitivity of discriminatory perceptual capacities. As such, it would appear that what is sensed can change with learning and experience. After all, if I discriminate based on what I perceive, then when I discriminate more finely, it would follow that I perceive more finely, as well.

This is the basic position of the New Look School of psychology. New Look holds that since what we know impacts our perceptual judgments and discriminatory abilities, then cognition affects the qualitative character of perception. Proponents of this position claim that perceptual processing is a kind of inference-making. The perceptual system, they claim, constructs a theory about the visual world and fashions qualitative character in light of that theory. [3]

The problem, however, is that changes in perceptual experience are compatible with cognitive impenetrability. This is because it is compatible with cognitive impenetrability that the

[3] See Jerome S. Bruner and Ann Leigh Mintum, "On Perceptual Readiness," *Psychological Review* 64 (1957): 123-52; Jerome S. Bruner and Ann Leigh Mintum, "Perceptual Identification and Perceptual Identification," *Journal of General Psychology* 53 (1955): 21-28; Jerome S. Bruner and Leo Postman, "On the Perception of Incongruity: A Paradigm," *Journal of Personality* 18 (1949): 206-23; Jerome S. Bruner and Mary C. Potter, "Interference in Visual Recognition," *Science* 144 (1964): 424-5.

intentional aspect of our perceptual experience is impacted by belief and knowledge. The examples that New Look employs do not prove that changes in perception happen on the qualitative level rather than at the level of judgment. After all, it may very well be the case that perceptual learning occurs only at the conscious level while the nonconscious sensory qualities of which we become conscious as a result of training, remain identical.

In the philosophical literature, this problem is captured by Wittgenstein's distinction between "seeing" and "seeing as".[4] When we are dealing with cognitive penetration, we must determine that the "seeing" and not just the "seeing as" is impacted by beliefs, knowledge, and experience. The problem, however, is that being able to distinguish "seeing" from "seeing as" does not entail that "seeing" is not impacted by cognition. Just like cognitive penetration does not follow as a result of perceptual learning, an immunity on the part of perceptual processing to thoughts or beliefs of a certain kind does not entail cognitive impenetrability.

The problem with arguments that attempt to reinforce the distinction between "seeing" and "seeing as" as one between pure qualitative sensation and intentional thought is that they often take cognition to be constituted by an untenably limited set of states and events.

For instance, Dretske, in considering the distinction between "seeing" and "seeing as" commits this classic mistake.[5] Dretske argues in the following fashion:

> If concepts were required for experience then I should not be able to have an experience without applying the corresponding concepts. I can have an experience of e.g. a fiddle without also believing of X that it is a fiddle. Therefore conceptual thinking is not required for experience.

But this argument is invalid. It does not follow from the fact that I can experience a fiddle, i.e., hear it or see it, without believing that I am hearing or seeing a fiddle, that I am not deploying *any* cognitive/conceptual apparatus whatsoever in my experience of seeing or hearing

[4] Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell Publishers, 1953), 165-94.
[5] Fred Dretske, *Naturalizing the Mind* (Cambridge, MA: MIT Press, 1995), 9-22.

that fiddle. That it is possible to experience objects and/or events while lacking a particular level of conceptual sophistication does not mean that no cognitive sophistication is necessary for the experience.

Relatedly, it is important to notice that Wittgenstein's own account of seeing does not commit him to a pure, raw perceptual given. For Wittgenstein, there is a difference between "seeing" and "seeing as," which means that not all seeing is an instance of interpreting. However, from this it does not follow that seeing is free from concepts and cognition. After all, if we look at Wittgenstein's examples of seeing it is clear that they are not confined to raw sensory inputs. Wittgenstein states that one *sees* a fork[6]—one does not see a metal pronged object *as a* fork.[7] To see a fork, however, clearly requires a high degree of cognitive sophistication.

My basic point here is that arguments that attempt to strictly distinguish the sensory from the cognitive are often as hasty as those that try to conclude that sensations are the products of thought by relying on instances of perceptual learning. Cognitive impenetrability theorists often assume that the types of things that qualify as cognitive or conceptual are only of the kind that human adults with linguistic mastery can employ. But the fact is that it would be much harder to argue that qualitative experience does not require cognition if we did not limit cognition to complex concepts such as fiddle-ness, but considered the comprehension of more basic facts and skills such as the mind-independence of objects or the ability to distinguish between music and noise, to be a cognitive achievement. This is not an arbitrary suggestion, for it is far from obvious that the comprehension of such basic facts is not contingent upon a set of cognitive capacities.

---

[6] Wittgenstein, 166.
[7] One can, however, see a cube *as* a box for a beetle (Wittgenstein, 165).

This problem has deep roots. The philosophical conception of cognition has always been decidedly narrow. From Aristotle to Hume to Fodor, cognitive states have been largely identified with propositional states and cognitive processes have been limited to the performance of truth-functional logical operations over these propositions. This notion of cognition, however, rules out nonlinguistic animals and pre-linguistic children from participation in the realm of the cognitive. Additionally, it fails to explain how anyone can know how to implement propositional knowledge in practice. This is because knowing how to do something is not reducible to knowing that *w* is the way to do it. When I explore the notion of skill as a process that is characterized by cognition, I will call into question this classical notion of intelligence. For now, I would simply like to emphasize the fact that answering the question about whether or not perception is cognitively penetrable depends, in part, on which states and events qualify as cognitive.

In addition to theoretical issues, there is also relevant empirical evidence to consider when thinking about cognitive penetration. For example, blind patients who undergo operations that restore their vision often report experiencing experiential blindness. These patients report experiencing a confusing cacophony of visual sensations after regaining their sight. Oliver Sacks reports that one such patient

> seemed to be staring blankly, bewildered, without focusing, at the surgeon, who stood before him, still holding the bandages. Only when the surgeon spoke—saying "Well?"—did a look of recognition cross Virgil's face. Virgil told me later that in this first moment he had no idea what he was seeing. There was light, there was movement, there was color, all mixed up, all meaningless, a blur. Then out of the blur came a voice that said, "Well?" Then, and only then, he said, did he finally realized that this chaos of light and shadow was a face—and indeed, the face of his surgeon.[8]

[8] Oliver Sacks, *An Anthropologist on Mars: Seven Paradoxical Tales,* (New York: Knopf, 1995), 114.

Such cases suggest that visual sensation is only comprehensible in light of the background knowledge one possesses about one's environment. In the above example, the recognition of the doctor's voice allows the patient to make sense of the jumble of visual properties that he is experiencing. Out of the confusion of light and color, he is able to form a visual representation of a face. But this "making sense of" by the perceptual system, it seems, is an instance of cognitive penetration. After all, the visual appearance of a face, rather than of a mess of sensations, qualifies as a change in qualitative character. That is, the arrangement and relationship of visual properties is a qualitative matter.[9] Compare Virgil's booming, buzzing, mess with our ordinary experience; even our nonconscious perceptions are of a fixed, ordered, and comprehensible world. So, it would seem that Noë is right when he says that such a "surgery restores visual *sensation*, at least to a significant degree, but… it does not restore sight."[10]

Further, experiments with inverting glasses suggest that skill in navigating one's environment is crucial to constructing a visual representation of that environment. Surprisingly, when subjects put on inverting lenses, there is not "an inversion of the content of experience (an inversion of what is seen) but rather a partial disruption of seeing itself."[11] And it is only after subjects learn to move around their spaces and develop appropriate motor responses that they regain the ability to see in an orderly fashion. This experiment indicates that the qualitative character of a perceptual experience is intimately related to spatial understanding and motor skill.

---

[9] We cannot say that sensation is so far down that cognitive impenetrability becomes trivial. " [I]t is crucial for our thesis that the notion of early vision includes, among other things, such functions as the individuation of objects, as well as the computation of what are referred to as "surface layouts"—the shape of the visible surface in our field of view. If it were not for the fact that early vision includes such complex features of the visible world, the independence theses would be trivial; everyone believes that *something* is detected by the visual system without regard for beliefs and expectations—otherwise we would only see what we wished to see" (italics in original). Zenon Pylyshyn, *Seeing and Visualizing* (Cambridge, MA: MIT Press, 2006), 50-1

[10] Noë, *Action in Perception,* 5.

[11] Noë, *Action in Perception,* 8.

After all, if meaningful visual experience is available only when one has the ability to interact with one's environment, then it appears that understanding and expectations impact the qualitative character of perception.

Finally, the anatomical structure of the brain suggests the possibility of cognitive penetration. The fact is that not only are neural structures "plastic," but many "top-down" connections exist between higher cortical regions and lower-order perceptual processing centers.[12] Both of these facts indicate the possibility of cognitive penetration.

Neural plasticity refers to the fact that while the brain is structured and ordered, it is not fixed or inflexible. Rather, the brain's neural connections have the capacity to adapt and change. And flexibility in neural processing is a prerequisite for cognitive penetration; fixed neural patterns simply could not accommodate the possibility of cognitive penetration. Further, it seems that the only reason that the brain should be plastic is so that it can improve and adapt when necessary; that is, so it can respond to knowledge, learning and experience. Clearly, if perceptual processing is cognitively penetrable, then qualitative character could be processed in ways that were more efficient relative to a creature's particular needs. So, neural plasticity supports the possibility of cognitive penetration though, admittedly, it does not entail it.

Further, although early perceptual processing is defined in a functional way, there is consensus on the fact that such functional processing is implemented in particular anatomical regions of the brain, namely the primary sensory cortices. And these areas, e.g., the primary visual cortex, the primary auditory cortex, the primary sensorimotor cortex, etc. all appear to have connections that run both from the bottom-up and from the top-down; e.g., information

[12] See Jean Buller, "Visual Perception is Too Fast to be Impenetrable to Cognition," in "Open Peer Commentary: Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," *Behavioral and Brain Sciences* 22 (1999): 370.

travels not only from the primary visual cortex to higher order cognitive regions, but also from higher cortical regions to early perceptual processing centers.[13] Evidence of these neural connections suggests that cognitive processing is wired to impact perceptual processing.

Additionally, studies have found that training can induce neural changes in primary sensory areas.[14] That is, studies have correlated skill-learning with plasticity in the primary visual, motor and auditory cortices. This is significant since it shows that areas that are functionally responsible for processing the qualitative character of a sensory event can be affected by learning and experience. However, it is important to note that changes in early sensory processing areas, which are correlated with training, do not guarantee cognitive penetration. That is, to demonstrate the cognitive penetration of perception, it is required that plasticity in the primary sensory cortices is (1) caused by cognitive impacts and not just adaptation, conditioning or exposure and (2) that such changes are not simply the result of intra-cortical reorganization.

While alone, none of this evidence amounts to a smoking gun, it does amount to a good reason to pursue theoretical clarification of our notion of cognitive penetration.

---

[13] Yasushi Miyashita and Toshihiro Hayashi, "Neural Representation of Visual Objects: Encoding and Top-down Activation," *Current Opinion in Neurobiology*, 10 (2000): 187-194 and Wu Li, Valentin Piëch, and Charles D. Gilbert, "Perceptual Learning and Top-Down Influences in Primary Visual Cortex," *Nature Neuroscience* 7 (2004): 651-57.

[14] Gilles Pourtois, Karsten S. Rauss, Patrik Vuilleumier, and Sophie Schwartz, "Effects of Perceptual Learning on Primary Visual Cortex Activity in Humans," *Vision Research* 48 (2008): 55-62l; Jonathan Fritz, Shihab Shamma, Mounya Elhilali, and David Klein, "Rapid Task-Related Plasticity of Spectrotemporal Receptive Fields in Primary Auditory Cortex," *Nature Neuroscience* 6 (2003):1216-223; Jerome N. Sanes, and John P. Donogue, "Plasticity and Primary Motor Cortex," *Annual Review of Neuroscience* 23 (2000): 393-415; G.H. Recanzone, C.E. Schreiner and M.M. Merzenich, "Plasticity in the Frequency Representation of Primary Auditory Cortex Following Discrimination Training in Adult Owl Monkeys," *The Journal of Neuroscience* 13 (1993): 87-103; Steven W. Cheung, Srikantan S. Nagarajan, Christoph E. Schreiner, Purvis H. Bedenbaugh and Andrew Wong, "Plasticity in Primary Auditory Cortex of Monkeys with Altered Vocal Production," *The Journal of Neuroscience* 25 (2005): 2490-2503; Christo Pantev, Bernhard Ross, Takako Fujioka, Laurel J. Trainor, Michael Schulte and Mitthias Schulz, "Music and Learning-Induced Cortical Plasticity," in *The Neurosciences and Music*, ed. Giuliano Avanzini, Carmine Faienza, Diego Minciacchi, Luisa Lopez and Maria Majno, 438-450 (New York: New York Academy of Sciences, 2003); Shaowen Bao, Edward F. Chang, Jennifer Woods and Michael M. Merzenich, "Temporal Plasticity in the Primary Auditory Cortex Induced by Operant Perceptual Learning," *Nature Neuroscience* 7 (2004): 974-981.

Now, most discussions concerning the cognitive penetrability of perception focus on whether the right part of perception is affected by cognition. These treatments focus on whether it is really the qualitative character of a perceptual state that is impacted by cognition, or just the intentional content.[15] For my purposes, this is an important but secondary question. Opponents of cognitive penetration accept that some changes in early perceptual processing do result from long-term learning and experience.[16] They do not deny that these changes occur, but rather, they deny that such changes constitute instances of cognitive penetration. In order to show that these changes are indeed instances of cognitive penetration, I construct a strong theoretical framework by which to approach skill and its connection to perceptual processing.

To reiterate, my main argument is this: *If* diachronic cognitive penetration is caused by skill, then these changes in perceptual processing constitute legitimate instances of cognitive penetration. As such, perceptual processing is not informationally encapsulated. I attempt to support this claim by: (1) presenting a detailed analysis of the definition of cognitive penetration, (2) arguing that propositional knowledge cannot account for practical know how, and (3) providing a definition of skill that highlights the particular aspect of skill that is cognitive in nature. Taken together, these considerations comprise an argument for the possibility of a legitimate instance of cognitive penetration resulting from the regular instantiation of skill.

2.1 Chapter outline: chapter two

I begin chapter two by considering Jerry Fodor's position concerning diachronic cognitive penetration. Fodor argues that long-term cognitive penetration does not constitute a legitimate instance of cognitive penetration. Because Fodor's argument is less an argument and

[15] See Jerry Fodor, "A Reply to Churchland's 'Perceptual Plasticity and Theoretical Neutrality'," *Philosophy of Science* 55 (1988): 188-98; Athanassio Raftopoulos, "Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception," *Cognitive Science* 25 (2001): 423-451.
[16] See Jerry Fodor, *The Modularity of Mind: An Essay on Faculty Psychology* (Cambridge: MIT Press, 1983), 82; Zenon Pylyshyn, *Seeing and Visualizing*, 87-8.

more an assertion, I attribute to him three plausible reasons for denying that diachronic cognitive penetration is actually an instance cognitive penetration:

> (a) The connections internalized into early perceptual modules as a result of long-term cognitive penetration are not propositional and, therefore, not cognitive.
> (b) Those states that are responsible for the long-term cognitive penetration of early perceptual systems are not propositional and, therefore, do not constitute an instance of cognitive penetration.
> (c) Early perceptual systems are automatic and, therefore, not cognitive.

I respond to these objections by analyzing the definition of cognitive penetration. Zenon Pylyshyn, an ally of Fodor's, defines cognitive penetration in the following way:

> If a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, that is it can be altered in a way that bears some logical relation to what a person knows.[17]

This definition elicits three major questions of its own:

(1) What part of perception must be affected by cognition in order for it to constitute an instance cognitive penetration?
(2) What kind of states are goals, beliefs and knowledge?
(3) What constitutes a semantically coherent or logical relationship to functional processing?

The answer to question (1) should be clear from the above discussion: the part of perceptual processing that is responsible for producing qualitative properties must be affected by cognition in order for it to qualify as a legitimate instance of cognitive penetration. I will assume that the right stage of perceptual processing (i.e., early perception) is affected by diachronic cognitive penetration, since Fodor and Pylyshyn admit this much.

In order to answer question (2), I explore the nature of intentional states. I begin by elucidating Fodor's Language of Thought (LOT) hypothesis. I argue that Fodor's commitment to the propositionality and conceptuality of thought undermines his own explanation of concept possession, which appeals to some highly cognitive-seeming things, like learning stereotypes. In

[17] Zenon Pylyshyn, "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," *Behavioral and Brain Sciences* 22 (1999): 343.

response to the inadequacy of LOT, I present several alternative theories of intentional content. I appeal to P.F. Strawson, Michael Dummett, Ruth Millikan, Susan Hurley, and Jose Bermudez's accounts of intentionality here. These theorists present accounts, which suggest that intentional content is not exclusively conceptual or propositional in nature.

The above considerations serve as both a partial response to Fodor's objection (b) and also as motivation to reconsider, which states or events qualify as cognitive. This fuels my subsequent inquiry into the cognitive nature of skills.

To respond to Fodor's objection (a), I explore which conditions must hold in order for cognition to impact the internal connections of a perceptual processing system. I argue that in order to be cognitively penetrated, intramodular connections need not be cognitive themselves but they do need to be related to cognition in the right way. So, simply because the connections internal to a perceptual system may not be propositional, does not preclude their potential for being cognitively penetrated. As such, Fodor's first objection fails.

My treatment of question (3) will lead me to ask what it means for a content to be semantically coherent or logically related to perceptual processing. In order to answer this question, I appeal to Fred Dretske's theory of minimal rationality. I apply Dretske's standards for minimally rational behavior to perceptual processing's relationship to cognition. What follows is that in order to qualify as an instance of cognitive penetration, cognitive content must be *explanatorily relevant* in accounting for the function of a perceptual system.

I end chapter two by responding to Fodor's objection (c). I take issue with the assumption that automatic systems are necessarily noncognitive. I do this by both considering cognitive capacities that are automatically instantiated and also by distinguishing between the level of rule implementation and the level of information processing. I claim that it is possible for a system to

automatically and inflexibly process information at one level according to rules that are variant and flexible at another.

2.2. Chapter outline: chapter three

In chapter three, I argue that the cognition characteristic of know-how[18] is not reducible to the intelligence of the propositional states that are often associated with skillful activity. Though skills are intimately connected to beliefs, desires, and knowledge, skill cannot be accounted for by appeal to intentional states. Following Gilbert Ryle, I argue that while propositional states certainly qualify as cognitive, the application of intelligence in action manifests a distinct, but equally legitimate, cognitive category.

I make this argument by exploring and responding to Jason Stanley and Timothy Williamson's objection to Ryle's distinction between knowing-how and knowing-that. I argue that Stanley and Williamson's objection to Ryle reduces knowing-how to an unintentional activity and, as such, overshoots its goal. Further, following Noë, I argue that Stanley and Williamson's own account does not solve the problem of knowing-how but merely relocates it.

I continue chapter three by arguing that propositional knowledge is neither necessary nor sufficient for knowing-how. In order to clarify this dissociation, I offer a thought experiment. I consider Bela Karoli and Mary Lou Reton, who are both presently unable to perform a standing layout on beam. I claim that while Mary Lou knows how to do a standing layout on beam, Bela, does not. Mary Lou knows-how because she has learned how and if she had the opportunity to perform a standing layout on beam under normal conditions, then she would be able to perform it successfully. Bela, however, could not perform the skill even under opportune circumstances. He knows *about* the skill, but he does not *know how* to perform the skill. In this example, Bela has more propositional knowledge than Mary Lou, but he does not have skill. Mary Lou has

[18] I use "know-how," "intelligent ability," and "skill" interchangeably.

skill but less propositional knowledge than Bela. Neither of them have the ability to perform the skill presently.

This example proves that propositional knowledge is not sufficient for knowing-how, since Bela has propositional knowledge but not skill. It also indicates that propositional knowledge is not responsible for successful skill instantiation. This is the case since, if propositional knowledge could account for knowing-how to *a*, then one should expect that the more facts one knows about how to *a*, the better one would be at *a*-ing.

Further, this example highlights the difference between knowing how to *a* and knowing that *w* is a way to *a* without locating the distinction in one's ability to *a* now. Thus, inability to perform a skill does not conflate distinctions in the knowledge bases of Bela and Mary Lou.

I continue chapter three by considering whether propositional knowledge is *ever* sufficient for knowledge-how. I conclude that if it is, such a case would be very much the exception. In this discussion, I consider John Bengson and Mark Moffet's claim that skill is conceptual. Bengson and Moffet hold that to possess concepts, one must know how to do certain things with them. Unfortunately, this argument alters the definition of a concept so that it includes that very part of knowing-how for which propositional content cannot account. As such, knowing-how has not been shown to be subject to the intellectualist legend.

Lastly, I consider what know-how is not. I claim that knowing-how is not the result of practical or theoretical reasoning, it is not propositional or conceptual, and it is not aimed at truth. I claim that because skill requires a high degree of sensitivity to the very embedded and particular properties of one's action-space, it cannot satisfy the criteria for concept-hood or propositionality and also, it cannot be the kind of thing that results from practical reasoning.

Also, I claim that know-how is aimed at success and not at truth. So, while knowledge-how is normative in that it can miss or hit its mark, its mark is not truth.

2.3 Chapter outline: Chapters four and five

In chapter four, I present and discuss three criteria for skill. They are:

(1) If *a* is a skill of *S*'s, then when *S* instantiates *a*, S must have some reason *r* for *a*-ing, and *r* must be responsible for *S*'s instantiation of *a*.
2) If *a* is a skill of *S*'s, then under conditions *c*, *S* is usually able to instantiate *a* successfully.
3) If *a* is a skill of *S*'s at time *t*, then there was some time *t*-1 that *S* could not instantiate *a* successfully, and *S* learned to instantiate *a* successfully as a result of practice.

Skills, I claim, are procedures that are learned and instantiated in particular ways. The first criterion of skill requires that skill instantiations are intentional actions. I define intentional actions as actions that are done for reasons, where a reason is composed of a pro attitude and an instrumental belief. Furthermore, for an action to be done for reasons, a reason must act as the cause of the action. In order to avoid concerns about the causal impotence of reasons, I encourage a functional reading of "cause" in this context.

Although an action must be performed for reasons in order for it to qualify as a skill instantiation, it does not follow that skill instantiations must be individuated according to the reasons for which they are performed. Instead, I argue that skill instantiations should be individuated extensionally and in an ontologically conservative fashion. On my account, a skill succeeds in being done for reasons if the reason for which it is done matches any of the extensional descriptions of the action that is caused by it. Further, if there are various descriptions of actions that have identical extensional boundaries, then those actions are identical. Thus, while reasons are opaque, skills are not. From this it follows that it is possible for an individual to know-how to do something without knowing that she knows-how to do it.

While skills are not identical to their instantiations, I claim that we can appeal to descriptions of skill instantiations as a way in which to type skills. In the same way that universals are typed by way of their particular manifestations, I propose that skill types should be identified by way of their particular instantiations. Thus, while a skill is a procedure and an instantiation is an action, the type of procedure that is identified as a skill is categorized according to the kind of behavior that is manifest when that procedure is instantiated.

In chapter four, I also argue that reasons need not be conscious or propositional in order to count as reasons. I claim that consciousness and intentionality are independent properties and, as such, the absence of one does not entail the absence of the other. Also, I argue that pro attitudes and instrumental beliefs can be accounted for by a Strawson-type feature placing language, a Gibsonian affordance account, or by the nonconceptual reasoning accounts of Hurley or José Bermudez. I present these alternative theories of intentional content in some detail in chapter two.

Turning to the second criterion of skill, I argue that success, or the ability to instantiate *a* successfully, is a necessary condition of skill. That is, if *a* is a skill of *S*'s, then *S* must be able to instantiate *a* successfully. Following Katherine Hawley, however, I argue that ability does not require success under any and all conditions, but instead, only under certain normal counterfactual conditions. So, to have the skill of reading, I do not have to be able to read in the dark without my reading glasses, but I do need to be able to read under circumstances that are normal for reading.

Further, one need not be successful 100% of the time in order to possess a skill. One should, however, succeed at *a*-ing reliably, relative to a skill context. I argue, again following Hawley, that canonical standards for success are not built into the definition of a skill. Rather, I

claim, that success standards are relative to situation and agent. This does not entail that standards of success are altogether subjective, but rather that they are subject to contextual constraints. I also argue that the language of counterfactuals is sufficient to sustain a robust notion of skill without appeal to dispositions.

Lastly, in chapter five, I argue that learning through practice guarantees that a skill is characterized by practical intelligence. I define practice as actions that are repeated for the sake of learning and that lead to the increased probability of success. While the intention to improve is required for an action to qualify as practice, I emphasize that the particular goal of the intention need not match the skill learned as a result of that practice. For instance, I can perform a tumbling routine every other day with the intention of improving my tumbling skills, and as a result of this repetition, I may also develop the ability to jump really high. Though I do not intend to learn to jump really high when performing my tumbling routine, because my actions are repeated with the intention to improve *at something*, my ability to jump really high counts as an ability that is learned as a result of practice.

Further, I claim that the method of skill acquisition is indicative of the nature of the thing that is learned as a result of that method. So, the types of methods that are used to learn a skill can help us to understand the nature of the things that are learned as a result of them. As such, the necessity of practice in developing a skill suggests that what is learned through practice is neither a propositional fact nor an intention. This is because if practice developed something other than an embodied, practical kind of intelligence, we would have no way of accounting for why it is that one needs to practice in order to learn a skill.

It is important to recognize that practical intelligence is necessary to account, not only for the success of a skill, but also for the particular manner in which a successful skill is instantiated

i.e., with precision, grace, elegance, force, etc. As such, I argue that practice ensures the practical intelligence of a skill, because through practice one develops a sensitivity to the relevant features of an action-space and the ability to respond to one's environment in a controlled and deliberate way. These capacities form the cognitive core of a skill. Further, empirical research in selective attention supports the idea that practice teaches one how to pay attention to one's environment. And studies in learning indicate that one can differentiate between the practical intelligence of a skill and the purely physiological impacts of repeated action.

Finally, I argue that even if skills become automatic as a result of practice, this does not undermine their practical intelligence. This is because there is a difference between a skill as a procedure that is learned and instantiated in a particular way and the particular instantiation of that procedure. Indeed, as I argue in chapter two, there is a distinction between the level of rule implementation and the level of information processing. This means that the level of information processing (skill instantiation) may be automatic, while the rules that the instantiation follows (the skill) are sensitive and responsive in ways that reflect intelligence.

I end this chapter by claiming that if the manner in which a skill is performed is related in an explanatorily relevant way to perceptual processing, then this relationship constitutes an instance of cognitive penetration. That is, if the practical intelligence of a skill is required to explain the functional processing of an early perceptual system, then that system is penetrated by cognition. I do not claim that diachronic cognitive penetration does in fact occur, but rather I provide the theoretical framework, which can be filled in with empirical evidence, in order to determine whether an instance of cognitive penetration has obtained.

CHAPTER 2

The goal of this chapter is to present a detailed analysis of the notion of cognitive penetration. This analysis will lay the groundwork for a response to Jerry Fodor's claim that long-term cognitive penetration does not qualify as a legitimate instance of cognitive penetration. This analysis will also produce a framework from which to argue that early perceptual systems are not modular. Significantly, I will not promote the idea that propositional thought may affect perceptual processing, but rather, that the instantiation of practical intelligence in action may impact the function that early perceptual systems compute.

1. Fodor, modularity, and long-term cognitive penetration

In *The Modularity of Mind*, Jerry Fodor argues that input systems, i.e., those early-perceptual systems that are responsible for processing sensory stimuli and producing the qualitative character of a perceptual event, are modular. By modular, Fodor means that these systems are characterized by the following properties: domain specificity; mandatoriness; limited access by central functions; fast; informationally encapsulated; productive of 'shallow' outputs; exhibiting characteristic breakdown patterns and exhibiting specific pace and sequence. For my purposes, I will focus on the claim that the early-perceptual systems are informationally encapsulated. It is this claim, which entails that early-perceptual input systems are cognitively impenetrable.

1.1 Modularity, informational encapsulation, innateness and functional individuation

Before proceeding to my main argument, I'd like to pause briefly to explore the relationship between modularity, informational encapsulation, innateness, and functionally

individuated systems. The term "module" was first coined by Noam Chomsky to refer to innate bodies of knowledge, which he posits as necessary to account for human language acquisition.[19] This term was then co-opted by Fodor to refer to the computational processing of functionally individuated systems that are characterized by the properties listed above. It is only modularity a lá Fodor that requires that the functional processing of modular systems are immune to the impacts of cognition. Fodor writes that,

> [T]he putative connection between Chomsky's kind of modules and Fodor's kind…is that bodies of innate knowledge are typically processed by encapsulated cognitive mechanisms; and vice versa, that encapsulated cognitive mechanisms are typically dedicated to processing of innate databases (e.g. the integration of innate information with sensory inputs early in the course of perceptual analysis). The idea that this is the typical relation between Chomsky's sort of modules and my sort continues to strike me as plausible.[20]

It is important to notice that informational encapsulation a lá Fodor and modularity a lá Chomsky are conceptually independent in several ways. That is, even if it turns out to be the case that there are innate bundles of knowledge with which human infants are born, their processing need not be informationally encapsulated. In principle, there is no reason to prevent innate knowledge from being supplemented, altered and augmented by a posteriori knowledge. Furthermore, there is nothing logically contradictory about the idea that innate knowledge is processed by structures that are sensitive to experiential learning and expertise. That it strikes Fodor as plausible that innate bodies of knowledge are processed by informationally encapsulated mechanisms is clearly not an argument for why those of us who do not find this idea plausible, must accept it anyway. That is, simply because we are born with systems that are

[19] See Noam Chomsky, *Aspects of the Theory of Syntax* (Cambridge: MIT Press, 1965).
[20] Jerry Fodor, *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology* (Cambridge: MIT Press, 2000), 57.

wired in particular ways does not entail that those systems are hard-wired, in the sense that their processing is immune to adaptation, alteration and change.

It is perfectly reasonable to suppose that particular functionally individuated systems may exhibit degrees of cognitive penetration. If a modular system, in the psychological sense of the term,[21] is not tossed around by any and every cognitive whim, this does not mean that such a system is not responsive, in certain limited and systematic ways, to the impacts of cognition. A slow learner still learns.

1.2 Jerry Fodor's informational encapsulation and long-term cognitive penetration

According to Fodor, the informational encapsulation of the early perceptual systems necessitates that input systems are immune to cognitive influences of any kind. That is, the function that an early perceptual system computes will remain identical regardless of the knowledge, skills and experience that an organism acquires. Importantly, this means that the perceptual properties, which result from early perceptual processing, are neither dependent on, nor sensitive to, what a creature knows or believes. Informational encapsulation requires that the qualitative character of a perceptual event is immune to the impacts of cognition. This must be the case since cognitively impenetrable early perceptual systems process sensory qualities and it is these properties that constitute qualitative character.[22]

Since Fodor takes informational encapsulation to be central to his theory of modularity,[23] from hereon, I will assume that the term "modularity" contains the notion of informational

[21] By this I mean a system that can be easily identified and isolated in functional terms. Fodor acknowledges that this is a watered-down notion of modulariy, where modularity just means that a process can be "put in a box." But, it is still crucial to recognize that cognitive penetration does not undermine the possibility of multiple, largely independent processing systems. If cognitive penetration is possible, it does not turn the mind into a whirling, whizzing mess of unindividuated activity.

[22] "It would not be a great exaggeration to say that early vision—the part of visual processing that is prior to access to general knowledge—computes just about everything that might be called a "visual appearance"" Zenon Pylyshyn, *Seeing and Visualizing* (Cambridge: MIT Press, 2003), 51.

[23] Fodor, *The Mind Doesn't Work that Way*, 58.

encapsulation. That is, if a system is cognitively penetrable (i.e., not informationally encapsulated) then that that system is not modular.

In order to demonstrate the informational encapsulation of early-perceptual systems, Fodor relies on, amongst other things, cases of persisting perceptual illusions. His claim is that if cognition were able to impact perception then what a creature knows or believes should change the way that things appear to that creature. But quite obviously, this does not occur. Take for example the Müller-Lyer illusion. Fodor reasons that if cognition could penetrate perception, then a person who *knows* that the two lines of the Müller-Lyer illusion are of equal lengths, should *see* those lines as equal in length. But she does not; that is the illusion. No matter how sincerely one believes that the two lines are of the same length, the line with the arrows pointed inward always appears longer. How then, could one claim that perception is penetrated by cognition?

This is a consideration that we cannot disregard if we are attempting to do justice to the relationship between cognition and perception. If we accept that the absence of change in conscious perceptual experience is indicative of a lack of change in qualitative character, as Fodor does and I am inclined to follow, then this is a strong counterexample to the possibility of cognitive penetration.[24] Cases such as this, and so many others like it, make it abundantly clear that if cognition is to impact perception, it certainly is not going to be a free-for-all. It will not turn out to be the case that every thought affects every perception.

[24] Of course, we should be careful to distinguish between the qualitative character of a perceptual event and the judgment of that qualitative character. One could claim that qualitative experience changes as a result of cognition, but that one's conscious experience does not reflect this change. Or, that qualitatively the lines appear similar in length but that we judge them as different, and our judgment is not impacted by our knowledge. But the latter claim is rather odd since there doesn't seem to be a dispute as to whether intentional states can affect judgment, and the former would simply imply that Fodor is wrong about modularity.

From the fact that some thoughts do not impact the qualitative character of a perceptual event, however, it does not follow that cognition, in general, does not impact the qualitative character of a perceptual event. Fodor is aware of this and so his argument for the informational encapsulation of the early-perceptual systems turns to the issue of long-term cognitive penetrability. Surprisingly, Fodor admits that cognition, in the long-term, can become internalized into the early perceptual modules. Fodor admits that cognition can impact the way an early perceptual system processes perceptual stimuli. However, he denies that such diachronic cognitive penetration counts as an instance of cognitive penetration. It is here, I think, that Fodor's argument gets into trouble.

Fodor states,

> [S]uch connection is not knowledge; it is not even judgment. It is simply the mechanism of contextual adjustment of response thresholds. Or, to put the matter metaphysically, the formation of interlexical connections buys the synchronic encapsulation of the language processor at the price of its cognitive penetrability *across time* (italics in original).[25] The information one has about how things are related in the world is inaccessible to modulate lexical access; that is what the encapsulation of the language processor implies. But one's experience of the relations of things is in the connections among lexical nodes.[26]

The problem, and this really is *the* problem, is that Fodor does not argue for why such automatic, internal connections cease to bear the right connection to cognition once they become automatic and internal. He simply states that this is the case. Clearly, Fodor must give us some reason to believe that the changes to automatic processing that result from cognitive penetration over time should not be considered instances of cognition impacting the qualitative character of a perceptual event, i.e., instances of cognitive penetration. Of course, there are a number of

[25] In this quote, Fodor is focusing on the lexical module, but he is clear that what will be true of language processing will also be true of perceptual input systems.
[26] Jerry Fodor, *The Modularity of Mind: An Essay on Faculty Psychology* (Cambridge: MIT Press, 1983), 82.

differences between paradigmatically cognitive events and their connection to perception and the potential diachronic affects of learning, experience and expertise on the processing of these input systems. However, Fodor does not produce any argument for why *these* differences amount to the difference between the cognitive penetration of perception and its denial. [27]

1.3 Detective work: deciphering Fodor's premises

Though Fodor does not explicitly argue for why long-term cognitive penetration is *not* a genuine case of cognitive penetration, I think that we can safely point to three assumptions, which may reasonably motivate his conclusion. The first two assumptions follow from Fodor's commitment to the Language of Thought hypothesis (from hereon, LOT). The third motivation, I postulate, follows from Fodor's commitment to the mandatory processing of modular systems.[28]

Famously, Fodor is responsible for articulating the very powerful LOT hypothesis. According to LOT, all thought is propositional in nature. A proposition is composed of a propositional attitude, such as a belief or a desire, that is directed at a propositional content, which is either a sentence in a real language or a sentence in the language of thought. One of the virtues of LOT is that it allows various attitudes to be directed at the same contents and various contents to be the objects of various attitudes. So, I can believe *that the weather will be cold tomorrow* and *that hot soup is great on cold nights* and I can fear *that the weather will be cold tomorrow* and you can believe and fear the same things, too.

---

[27] Pylyshyn confronts the same problem as Fodor. He says: "it is consistent with the present framework that new complex processes could become part of the early vision system over time: cognitive impenetrability and diachronic change are not incompatible" Pylyshyn, *Seeing and Visualizing*, 88. Like Fodor, why we should buy that modularity is sustained when long-term cognitive penetration occurs is simply stated, but it is not argued.

[28] It should be clear that Fodor is not necessarily committed to each of these objections (after all, I'm guessing). It is very likely, though, that he holds at least one. The others are worth discussing just in case others consider them plausible reasons to reject the possibility that diachronic cognitive penetration undermines the informational encapsulation of input systems.

The propositional content, or the sentence at which a propositional attitude is directed, is also internally structured. Such sentences are compositional, which means that they are made up of atomistic parts or concepts that can enter into meaningful relationships in an infinite number of sentences. It is important to note that these sentences are not necessarily in any real language. Rather, Fodor argues that thought is structurally, logically and grammatically like sentences in a natural language, except without necessarily being in language; hence, the language of thought. I will address compositionality and its commitments and implications further in section 2.4. For now, we should simply note that this fact about propositional thought is genuinely explanatory in accounting for the systematicity and productivity of thought and language.

Now, because Fodor is committed to this complex but narrow understanding of thought, it should be quite obvious that if a state is not propositional, then Fodor would not classify it as cognitive. After all, that's exactly what LOT says: for something to be an instance of thought, that something must be propositional in nature. Therefore, we can formulate two possible reasons for Fodor's conclusion that long-term cognitive penetration does not qualify as a legitimate instance of cognitive penetration: (1) The connections internalized into an early perceptual module as a result of long-term cognitive penetration are not propositional, and therefore not cognitive, and (2) Those states that are responsible for the long-term cognitive penetration of early perceptual systems are not propositional, and therefore they do not constitute instances of cognitive penetration.

The third motivation for Fodor's claim that long-term cognitive penetration is not an instance of cognitive penetration, may follow from his commitment to the automaticity or mandatoriness of input processing systems. After all, it is often noted that what differentiates cognitive systems from non-cognitive ones is that their responses possess a certain degree of

flexibility. They are not automatic like reflexes or instincts. So, Fodor may be relying on the fact that the connections internalized into early perceptual systems, as a result of long-term cognitive penetration, do not exhibit any of this characteristic flexibility in order to conclude that they do not constitute instances of cognitive penetration.

In the following sections, I will investigate whether any of these three objections legitimize the conclusion that long-term cognitive penetration is not an instance of cognitive penetration.

2. Cognitive penetration: definition and analysis

In order to assess the force of the objections that I've ascribed to Fodor, it will be useful to examine the definition of cognitive penetration. This will be essential, not only as a way to respond to Fodor, but also in order to ground the positive account of long-term cognitive penetration that I will present in the following chapters.

Zenon Pylyshyn, a proponent of the cognitive impenetrability of perception and an ally of Fodor's,[29] has defined cognitive penetration in the following way:

> If a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, that is it can be altered in a way that bears some logical relation to what a person knows.[30]

A thorough analysis of this definition will require that we are absolutely clear on three issues: (a) What part of a perceptual system must be affected by cognition in order to qualify as an instance of cognitive penetration? (b) What kind of states are goals, beliefs and knowledge such that they are representative of cognition? And (c) What constitutes a semantically coherent or logical relationship to the functional processing of a perceptual system?

---

[29] I mention this fact only to assure the reader that I am not using a definition of cognitive penetration that is biased in my favor in order to argue against Fodor on grounds that he would not accept. This dispute is substantive.

[30] Zenon Pylyshyn, "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," *The Behavioral and Brain Sciences* 22 (1999): 343.

Before answering these questions, I'd like to note that the above definition does not stipulate that changes in perception must immediately follow from a connection to occurrent thoughts. Such a limitation, of course, would rule out long-term cognitive penetration as a potential instance of cognitive penetration not by argument, but by fiat. This would not only beg the question, it would also, as Churchland points out, argue against a straw man.[31] The shear long-term-ness of cognitive penetration cannot be the reason that changes in perceptual processing over time fail to qualify as instances of cognitive penetration.

Of course, the idea that occurent thoughts or beliefs do not cause immediate changes in perceptual processing is a robust thesis in and of itself. But this is not the modularity thesis. That is, though proponents of modularity may be committed to the notion that only the impacts of occurrent, propositional thoughts should count as instances of cognitive penetration, this cannot exhaust our dealings with this issue. After all, we are asking not just if cognitive states, as Pylyshyn and Fodor define them, impact perceptual processing, but rather, if cognition affects perception, full stop. The thesis about propositional states and their relationship to perception is an important one, but it should not be confused with the larger question of the relationship between qualitative character and cognitive content.

2.1 Early perceptual processing

The answer to question (a), what part of perception must be affected by cognition in order for it to count as cognitive penetration, is uncontroversial. Everyone agrees that early perception, or that part of perception that is responsible for processing qualitative properties, must be affected by cognition in order for it to constitute an instance of cognitive penetration.

[31] "[I]f Fodor is attacking the view that perceptual processing always (or even usually) responds directly and immediately to changes in one's theoretical commitment, then he is attacking a straw man. This is not a view that anyone has defended" Paul Churchland, "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor," *Philosophy of Science* 55 (1988), 176.

Because both Fodor and Pylyshyn admit that there are changes in early perceptual processing that result from diachronic cognitive penetration, I will not be concerned to argue this point. From hereon, I will assume that the right part of perceptual processing is impacted by long-term cognitive penetration.

2.2 Cognition and intentionality

To address question (b), it is essential to identify which types of states Pylyshyn is referring to when he appeals to beliefs, goals and knowledge as representative of the category of cognition. After all, examples are not definitions and we must have some way to distinguish between states that will, like beliefs, goals and knowledge, turn out to be cognitive, and those states that will not.

First off, it is crucial to agree that Pylshyn appeals to beliefs, goals and knowledge as paradigmatic instances of cognitive states. This must be so, since what Pylyshyn is concerned to define is a particular relationship between cognition and perception. He is not, after all, presenting a definition of *propositional* penetration, but of *cognitive* penetration. So, in the back of our minds, we should beware of the possibility that the category of states that is best represented by beliefs, goals and knowledge may not exhaust the category of cognition. If it turns out that this is the case (and not to spoil the surprise, but it will) then we must amend Pylyshyn's definition to encompass cognitive states more generally.

Quite clearly, beliefs, goals and knowledge are examples of intentional states. They are the types of states that Franz Brentano identifies as having the property of being about or directed at something. [32] Contrast this with artifacts and states of affairs, which are not *about* anything, but just *are*.

[32]Franz Brentano, *Psychology from an Empirical Standpoint*, translated by A.C. Rancurello, D.B. Terrell, and L. McAlister (London: Routledge, 1973), 88.

Importantly, intentional states can be true or false. As opposed to artifacts, representational states can be veridical or mistaken because they are not just part of the world, but because they represent some part of the world. And once one represents, one is faced with the possibility of misrepresenting. Not accidentally, intentional states are those states that are best accounted for by LOT; they are the kinds of states that are most likely to be propositional and conceptual in nature.

Intentional states are most commonly attributed to adult language users and they are usually individuated according to the types of reports that persons can give about them. For instance, the best way to ascertain whether or not I possess the belief that the tree has a broken branch is to ask me if I believe the tree has a broken branch. My belief report will be in a natural language and my propositional attitude will usually be followed by a *that*-clause. For instance, if prompted, I would say, "I believe *that* the tree is at the top of the hill" or I would say "I hope *that* we get to the top of that hill soon" or "I know *that* the tree at the top of the hill has a broken branch."

Given that the clearest expression of intentional states are both in language and in propositional form, we should ask whether one or both of these properties are necessary for some state to qualify as intentional. Since we know that these properties can come apart and since Fodor does not require that propositional states are in language, I will forego presenting arguments in favor of the fact that intentionality is not necessarily a linguistic affair. Instead, I will investigate whether intentional states can be non-propositional.

2.3 The Relationship between cognition, intentionality and propositionality

In this section, it will be my goal to argue that cognitive states are not necessarily propositional in nature. I will do this by first clarifying the nature of propositional content and

next, by gesturing to some very cognitive seeming activities that do not obviously qualify as propositional. My intention here is not to present a full-fledged theory of non-propositional cognition, but rather, to sketch out the conceptual space that one must occupy in order to respond to Fodor's objection (2) to long-term cognitive penetration: Those states that are responsible for changes in early perceptual systems over a period of time are not propositional, and therefore do not constitute instances of cognitive penetration.

To begin, I'd like to say a word about nonconceptual content. Many philosophers have argued that there is an aspect of perception whose content is not conceptual.[33] What they usually have in mind is that perception is a hybrid state composed of both phenomenological and intentional content. Though it is clear that the intentional aspect of perceptual content is conceptual (for adult human beings, anyway) they claim that the "what it's like", or the phenomenal character of perception, is nonconceptual.

I should explicitly note that even though one may admit that perceptual states are partly constituted by nonconceptual content, as Fodor does when he speaks of them as iconic,[34] this does not thereby entail that one must accepts that *cognitive* states are non-propositional. This is because, for most philosophers, the nonconceptual character of perceptual states presupposes that this aspect of their content is not cognitive. We have moved no further in proving that cognition can be non-propositional by admitting that the phenomenal character of perceptual states may be

[33] See, for instance, Gareth Evans, *The Varieties of Reference* (Oxford: Oxford University Press, 1982); Fred Dretske, *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981); Tim Crane, "The Waterfall Illusion," *Analysis* 48 (1988): 142-147; Christopher Peacocke, "Analogue Content," *Proceedings of the Aristotelian Society* 60 (1986): 1-17; Christopher Peacocke, "Perceptual Content," in *Themes from Kaplan* ed. Joseph Almog, John Perry, and Howard Wettstein (New York: Oxford University Press, 1989); Christopher Peacocke, *A Study of Concepts* (Cambridge: MIT Press, 1992); Michael Tye, "Nonconceptual Content, Richness, and Fineness of Grain," in *Perceptual Experience*, ed. Tamar Szabo Gendler and John Hawthorne (Oxford: Oxford University Press, 1996); Michael Tye, *Ten Problems of Consciousness* (Cambridge: MIT Press, 1995); Sean Kelly, "Demonstrative Concepts and Experience," *The Philosophical Review* 110 (2001): 397-420.

[34] Jerry Fodor, *LOT 2: The Language of Thought Revisited* (Oxford: Oxford University Press, 2008), 171.

nonconceptual. Though, of course, this claim does commit one to the idea that representational content is not necessarily propositional in nature.

Fundamentally, however, this issue is orthogonal to my purposes in this chapter. Though it would certainly make my case stronger if, in fact, phenomenal character were conceptual,[35] such a condition is not required to allow for the possibility of cognitive penetration. It is not necessary that the phenomenal aspect of a perceptual state is conceptual in order for it to be cognitively penetrable, just so long as it bears the right relations to conceptual or cognitive content. For this reason, I will not be concerned to enter the debate regarding the conceptual or nonconceptual nature of qualitative content. If it turns out that phenomenal content is conceptual then perception, it seems, must be cognitively penetrable. Great. If we do not have to accept this position, as I'm inclined to think that we do not, perception can still be cognitively penetrable in virtue of being related to cognition in the right way. The latter is the claim that I will be interested in exploring.

My goal then, is not to argue that qualitative character is conceptual, but rather, to focus on the possibility that cognitive states may be nonpropositional. To better organize the landscape, it may be helpful to identify the possible positions that one may hold regarding the nature of, and relationships between, intentionality, cognition and propositionality:

(1) Intentional states exhaust the category of cognition (i.e., all and only intentional states are cognitive) and all cognitive states (ergo all intentional states) are constituted by propositional content.[36]

---

[35] As do intentionalists and representationalists such as David Armstrong, *A Materialist Theory of the Mind* (London: Routledge and Kegan Paul, 1968); Alex Byrne, "Intentionalism Defended," *Philosophical Review* 110 (2001): 199-239; and William Lycan, "In Defense of the Representational Theory of Qualia (Replies to Neander, Rey and Tye)," in *Philosophical Perspectives* 12 (1998): 479-87. See also, John McDowell, *Mind and World* (Cambridge, MA: Harvard University Press, 1994).

[36] This is Fodor's position.

(2) Intentional states exhaust the category of cognition (i.e., all and only intentional states are cognitive), and some cognitive states (ergo some intentional states) may be constituted by nonpropositional content.

(3) Intentional states do *not* exhaust the category of cognition, and while all intentional states are necessarily propositional, other cognitive states may be constituted by nonpropositional content.

(4) Intentional states do *not* exhaust the category of cognition, and both cognitive and intentional states may be constituted by nonpropositional content.

The last option would be (5) Intentional states do *not* exhaust the category of cognition, but both intentional and cognitive states are necessarily propositional. But it seems that this would collapse into option one, since it is unclear that these other propositional cognitive states could be anything but intentional. For the same reason, (6), Intentional states do *not* exhaust the category of cognition, and while intentional states may be constituted by nonpropositional content, cognitive states are necessarily propositional, is not a genuine option.

At minimum, if (2), (3) or (4) are true, then it may be possible for nonpropositional states to cause cognitive penetration. So, while it is not necessary that I prove that intentional states may be nonpropositional in order to show that other potentially nonpropositional cognitive states could be responsible for cognitive penetration, it will certainly be worth exploring this possibility here.

2.4 The nature of propositional content

In order to determine which states count as examples of nonpropositional cognition, it is a good idea to have a working definition of "propositional content." According to Fodor, propositional content is necessarily compositional. In *LOT Revisited*, Fodor writes,

> LOT 1 failed, just about entirely, to recognize the centrality of compositionality in constraining theories about the semantics of mental representation; that is, the implications of the requirement that the content of a thought is entirely determined by its structure together with the content of its constituent concepts.[37]

The compositionality of propositional thought is explanatorily potent in accounting for both the productivity and the systematicity of thought. And these are features that Fodor considers to be definitive of cognition. To reiterate, the compositionality of propositional content is at the very heart of the LOT hypothesis.

What compositionality requires is that cognitive states, such as beliefs and desires, are composed of concepts. Based on Fodor's account, there cannot be propositional content if there are no concepts, since intentional content is determined exclusively by concepts and their structure. As he states, "If you are going to have beliefs in your ontology, you are also going to have concepts, since the latter are the constituents of the former."[38] It should be clear then, that if some state is not conceptual, then it is not propositional either. And, according to Fodor, it is not intentional or cognitive, full stop.[39]

Accordingly, we should inquire into the nature of concepts. The first thing to notice is that on the LOT hypothesis, concepts are not only defined as, but must be defined as, atomistic. Individual concepts are independent of each other and also of their particular environment.[40] This means that, "in principle one might have any one concept without having any of the others."[41] This also means that concepts are not identified with any one particular context. It is this atomistic quality of concepts that accounts for their ability to appear in different contents and

---

[37] Fodor, *LOT 2: The Language of Thought Revisited*, 17.
[38] Fodor, *LOT 2: The Language of Thought Revisited*, 131.
[39] There is a way of using "cognitive" where it refers to any process that contributes to cognition, or any process that takes place in the brain. I hope it is clear, that it is not this weaker sense of the word that I am using here.
[40] It is interesting to note that this characteristic of concepts makes it impossible that concepts are definitions. This is because, quite obviously, definitions require concepts and so, any concepts would be dependent on others.
[41] Fodor, *LOT 2: The Language of Thought Revisited*, 141.

as the constituents of more complex concepts. That is, the ability of concepts to break free from their particular situation allows them to show up in many others.

This same feature allows us to draw a straight line between the belief "that there is a red house over there" and the belief "that there is a red car over there." Because the same Red concept appears in each of these sentences, we can easily see what the car and the house have in common. Notice how this would change, if the "red" of each sentence were necessarily tied to its environment. Obviously, the particular shade of red paint on the car and the particular shade of red paint on the house will differ. As a result, it would become quite problematic to explain why the same word is used to represent these variable features. If the concept Red could not be abstracted and reified away from its particular situation, that is, if it were not atomistic, then it would be quite difficult to explain what these two propositional contents share.

As Gareth Evans has explicated, concepts are both generalizable and recombinatorial. This means that for one to possess a concept, one must be able to think of that concept without applying it to any particular situation; one must be able to entertain a thought containing that concept without knowing whether it is true or false. Further, if one possesses a concept then one can apply it to various situations; one can think of that same concept in different contexts.

> It is a feature of the thought-content *that John is happy* that to grasp it requires distinguishable skills. In particular, it requires possession of the concept happiness—knowledge of what it is for a person to be happy; and that is something not tied to this or that particular person's happiness. There simply could not be a person who could entertain the thought that John is happy and the thought Harry is friendly, but who could not entertain—who was conceptually debarred from entertaining—the thought that John is friendly or Harry is happy.[42]

As we can see, what it is to be a concept is exactly what explains why concepts can build and recombine in such a way as to produce the variability and complexity of thought and language. It is the context-independent nature of concepts that explains how various contents

[42] Evans, 102-3.

can share constituent parts. If concepts did not break free of their environments, then one could not, as Fodor requires, bring them "before the mind as such."[43] One could not bring them or move them anywhere as such. They would be glued to their particular environment. They could not generalize or recombine. They'd be stuck.[44]

It is exactly this requirement that makes it so unlikely that nonlinguistic animals and pre-linguistic children, and perhaps surprisingly, even human adult doers and movers, are best explained by appeal to propositional thought. The bar for propositional thought is high and the fact remains that there is no need to posit concepts in order to explain all sorts of basic abilities and behaviors. In fact, many intelligent actions resist explanation in propositional or conceptual terms.

For example, it would seem that children at a very young age can desire particular sweet things (like that piece of chocolate cake) without being able to think of sweetness, "as such." And dogs can believe that their master is at the door, without having the concept of some other dog's "master." Further, as Ruth Millikan has argued,

> [W]hen hiking a steep path, you may take a brief moment to see the best way to get a leg up. Or before crossing a stream on scattered rocks, it may take you a while to see a good way to cross without wetting your feet. This sort of trying to see a way seems entirely different from practical inference as inference is usually described.[45]

In chapter 3, 4, and 5, I will cash out in some detail the nature of non-propositional, non-intentional cognition. For now, however, it is vital that we understand this one crucial fact: the atomistic requirement of concepts to move freely and recombine and generalize is not identical

---

[43] "[A] sufficient condition for having the concept C is: being able to think about something *as* a C (being able to bring the property C before the mind as such, as I'll sometimes put it)" Fodor, *LOT 2*, 138.

[44]John McDowell, *Mind and World* (Cambridge: Harvard University Press, 1994) has argue that concepts can be non-general, but I'm not sure that this isn't just changing the definition of concepts and then claiming that perception is conceptual. If we agree with Fodor on the nature of concepts, we don't need to even explore this notational grievance.

[45] Ruth Garrett Millikan, "Styles of Rationality," in *Rational Animals?* ed. Susan Hurley and Mathew Nudds, 120 (Oxford: Oxford University Press, 2006).

to a more basic ability to discriminate features of a perceptual array, recognize similarities and differences amongst features that one has previously encountered, or to group instances into a stereotype. This ought to be clear from Fodor's own explication of concept formation.

Fodor explicitly differentiates between stereotypes and concepts.[46] For Fodor, learning a stereotype is a necessary, but not a sufficient condition for acquiring a concept. To become a concept, a stereotype needs something more (some locking attractor stuff).

It should be fairly obvious, however, that in order to learn a stereotype a creature requires some pretty sophisticated cognitive capacities. After all, to form a stereotype, it is necessary that a creature have the ability to recognize, discriminate and group salient features of its environment appropriately. It must be able to detect good examples of the stereotype that it is forming; they must be paradigmatic instances of the type that the stereotype is typing. But, if an animal only has the ability to form stereotypes, or to place features, rather than to think conceptually in subject-predicate form, if it cannot reidentify objects, or bring them "before the mind as such," then that creature is not thinking. Again, for Fodor, if a creature does not have concepts, then that creature does not have thoughts. Thoughts are necessarily propositional and propositions are necessarily constituted by concepts. It follows then, that a creature that can learn stereotypes, but cannot or has not acquired concepts, is not exercising intelligence; that creature is not in possession of cognitive states.

Importantly, there are several options for exploring the nature of intentional thought, which do not require that cognitive states are necessarily constituted by concepts. For example, we can begin with an appeal to a Strawson-type feature-placing language. Like P.F. Strawson,[47] Adrian

[46] Fodor, *LOT 2*, 150.
[47] P.F. Strawson, *Individuals: An Essay in Descriptive Metaphysics* (New York: Anchor Books, 1959).

Cussins [48]and Austen Clark argue,[49] the capacity to recognize features in an environment does not imply that those features are bounded properties that can be thought of in the absence of the particular situation in which they occur. Those features need not be predicate qualities that can be identified and re-identified.

So, we can imagine some creature wanting to eat that animal over there, without attributing to that creature the concept of "hunger," "animal" or "over there," generally. Possessing such a primitive but contentful psychological state does not entail being able to identify it as the same state at different times, to be able to recognize it in others, or to think about it as such. It only requires being able to distinguish this state from other states, and this much can be done with features or sortals rather than concepts. That is, similarity and difference relations can account for a creature's ability to distinguish one content from another without also guaranteeing that the creature can recognize these contents as reidentifiable particulars. Discriminating between objects does not require discriminating between them as conceptualized particulars. After all, it is fairly easy to imagine a dog being able to differentiate between being hungry for that steak over there, but not the watermelon, without having concepts of either in any full-blown way.

Likewise, there are several explanations of practical reasoning, which avoid commitments to propositional thought or the capacity for full-blown inferential promiscuity. For instance, Ruth Millikan[50] and Michael Dummet[51] have, following J.J. Gibson,[52] proposed accounts of reasoning where affordances act as the basis for practical rationality and intentional action. Affordances are not the conclusions of logical inferences, but rather perceptually salient possibilities for

---

[48] Adrian Cussins, "Content, Embodiment and Objectivity: The Theory of Cognitive Trails," *Mind* 101 (1992): 651-88.
[49] Austen Clark, "Feature-Placing and Proto-Objects," *Philosophical Psychology* 17 (2004): 443-69.
[50] See Millikan, "Styles of Rationality," 117-126.
[51] Michael Dummett, *Origins of Analytic Philosophy* (USA: Harvard University Press, 1993).
[52] J.J. Gibson, *The Ecological Approach to Visual Perception* (Boston: Houghton Mifflin, 1979).

action. An affordance is an objective property that is relative to an agent and, when recognized, qualifies as a basic type of instrumental belief about what can be done given particular motivational forces. Affordances are instrumental beliefs that are directly perceived in an environment. Because affordances are particularly embedded in particular environments, they cannot satisfy the requirements for conceptual thought.

A slightly more sophisticated theory of non-conceptual practical rationality has been forwarded by Susan Hurley.[53] Hurley has argued that conditional, means-ends reasoning can exhibit degrees of flexibility and abstraction, without reaching the level of a fully conceptualized, truth-preserving, logical capacity.[54] That is, conditional reasoning may remain domain-specific, but still serve as the basis of rational action just as long as means and ends come apart and recombine in certain limited ways. This type of proto-reasoning can provide the basis for variant behavior that is done for reasons, but not fully conceptualized ones.[55]

A creature can reason about his actions and his environment by recognizing contingencies and consequences in a particular setting without necessarily being able to generalize this type of reasoning to all environments that are duplicate in logical structure.[56] Still, the ability to

---

[53] Susan Hurley, "Making Sense of Animals," in *Rational Animals?* ed. Susan Hurley and Matthew Nudds, 139-171 (Oxford: Oxford University Press, 2006).

[54] "[N]on-human animals can occupy islands of practical rationality: they can have domain-specific reasons for action even though they lack full conceptual abilities" Hurley, 139.

[55] "Perceptual information leads to no invariant response, but explains actions only in the context set by intentions an the constraints of at least primitive forms of practical rationality. Perceptions and intentions combine to make certain actions reasonable and appropriate from the animal's perspective. Means and end can decouple: an intentional agent can try, err, and try again, can try various different means to achieve the same end. The holism of intentional agency provides a minimal, coarse kind of recombinant structure: an intentional agent has the ability to combine a given intention with different perceptions, given ends with different means. Moreover, intentional agency is not merely a complex pattern of dispositions. It essentially involves normative constraints, relative to which mistakes make sense; actions can be inconsistent or instrumentally irrational" Hurley, 142.

[56] For example, a chimpanzee can recognize the rule "if I choose the smaller number then I will get more candy", without being able to generalize this rule to "if I choose the smaller amount of candy then I get more candy." The former can be an instrumental belief without having the capacity to be abstracted into a general rule and applied in any environment. See S.T. Boyson and G. Bernston, "Responses to Quantity: Perceptual Versus Cognitive Mechanisms in Chimpanzees (*Pan troglodytes),*" *Journal of Experimental Psychology and Animal Behavior Processes* 21 (1995): 82-86.

recognize certain available options and choose an appropriate response, even when the recognition and choice are limited and embedded in a particular set of circumstances, is still a form of instrumental reasoning. Reasons need not be fully environmentally independent in order for them to be reasons for action. As such, this basic means-ends reasoning can serve as the basis for nonpropositional beliefs and intentions.

Moving towards a greater degree of sophistication, José Bermúdez has proposed a kind of proto-logic that paves the way for attributing practical reason to non-linguistic creatures. This proto-logic is underpinned by a combination of a creature's ability to possess contrary concepts[57] and recognize environmental regularities. Bermúdez states,

> the ability to deploy pairs of contrary concepts (without an explicit understanding of the notion of contrariety) and to be sensitive to causal regularities in the distal environment can provide animals with the tools for relatively complex forms of practical reasoning—and certainly for forms of practical reasoning that are sufficiently complex to underwrite the extension of belief-desire psychology to non-linguistic creatures.[58]

His account gives us a way to think about animals that are not capable of linguistic, propositional, fully-abstract notions of contrariness or causation, as still possessing the capacity to reason instrumentally. That is, such creatures are capable of performing proto-logical reasoning, which closely resembles *modus tollens* and *modus ponens*. This sort of understanding of action and regular effect can underpin a proto-conceptual notion of intentionality and reasoning.

These considerations, I think, legitimately raise the possibility that not all cognitive states are propositional in character. This is certainly not a fully worked out defense of such a position, but it is enough to warrant doubt concerning the LOT hypothesis. What should be clear is that if

---

[57] These contrary concepts should not be thought of as full-blown reidentifiable particulars. José Bermúdez, *Thinking without Words* (Oxford: Oxford University Press, 2003), 129 calls these "primitive logical concepts."
[58] Bermúdez, 137.

cognitive states can be nonpropositional then the right relationship to the right sorts of nonpropositional states will qualify as an instance of cognitive penetration. As such, Fodor's second purported reason for rejecting cognitive penetration, that is, because the states that affect early perceptual processing are nonpropositional, is up for grabs.

The above considerations, I believe, provide cursory justification for the plausibility that cognitive states are much wider in scope, than LOT would have us believe. It will be the purpose of the following chapters to forward a positive account of non-propositional, non-intentional cognition.

3. Semantic Coherence and Logical Relations

3.1. Cognition is not a Transitive Property

In order to respond to Fodor's first objection to long-term cognitive penetration's constituting a legitimate instance of cognitive penetration, I think it would be useful to continue elucidating the commitments that follow from the definition of cognitive penetration. A thoughtful analysis of the definition of cognitive penetration should make it clear that there is an important distinction between a cognitive system and a cognitively penetrable system. Recall Pylyshyn's definition, which stated that for a system to qualify as cognitively penetrable, that system must bear the proper logical or semantic relationship to cognitive states or events. So, in order to ascertain whether a system is cognitively penetrable, the proper question is not: "is it a cognitive system?" or "does the system have intentional states?" but rather, "does the system bear the right relationship to cognitive states?" This is the appropriate question since, according to the definition of cognitive penetration, what is necessary for a system to be cognitively penetrable is not that it be a cognitive system in its own right, but rather, that the system exhibits a sensitivity or a responsiveness to cognitive states.

This much should be clear since a sensitivity or responsiveness to cognition does not thereby transform the sensitive system into a cognitive or intentional one. Another way to put this is: the relationship to cognition is not transitive. So, being affected by cognition does not entail becoming cognitive. And the same holds true for intentionality. That some system is responsive, in the right way, to beliefs, goals and knowledge does not thereby produce intentional states in that system.

An example will help to elucidate this point. Let us think of the motor system. The motor system is not a cognitive system, and yet it is responsive to cognitive states. That is, the function of the motor system is directly related to the goals, beliefs and knowledge of an agent. These connections are not simply causal, they are explanatory. That is, the content of an agent's goals and beliefs have a direct impact on the activity of the motor system. So, if an agent wants to pick up a book instead of a glass, the motor system will respond to these various desires in appropriate ways—by reaching and shaping one's hand accordingly. If one believes that the glass is hot, then one will move one's hand in a guarded fashion towards it. If one thinks that the book is heavy, one will begin to lift it with more force than if one believed that it was light. In this way, the activities of the motor system are sensitive to the content of the goals and beliefs of an agent.

Clearly, however, the motor system does not become a cognitive system simply in virtue of its sensitivity to cognitive states. Likewise, we should keep in mind that early perceptual systems do not need to become cognitive systems in order for them to be cognitively penetrable systems. They need not be cognitive or intentional, just so long as they are related to cognition in the right way.[59]

[59] I think that this is a fairly powerful but often overlooked fact about functional systems, perceptual or otherwise. I think that it is a confusion about this point that makes it tempting to say that perceptual systems make inferences, or

From this discussion, it should be clear that Fodor's first presumed objection to long-term cognitive penetration, that is, that the connections internalized into an early perceptual module as a result of long-term cognitive penetration are not propositional, and therefore not really instance of cognitive penetration, misses the mark. It is not required that the connections internal to an early perceptual system be cognitive in order for the system to be cognitively penetrable. This, I hope, is quite clear: a system can be sensitive or responsive to cognition without itself being a cognitive system.

3.2 Semantic coherence and logical relations

But exactly what is the right connection to a cognitive state such that this connection qualifies as an instance of cognitive penetration? How can we cash out the notions of semantic or logical coherence in such a way as to get the right type of relationship between a cognitive state and a functional system? This question will be vital not only in judging when instances of cognitive penetration occur, but also in constructing a positive account of diachronic cognitive penetration.

We certainly do not want just any logical relation to count as cognitive penetration. After all, if everything is cognitively penetrated, then the argument becomes trivial. But which connection should count as the right one? When I rearrange my plants, do my plants thereby become cognitively penetrated? When I diet and my stomach shrinks, is my stomach cognitively penetrated? These intentional activities and their results are certainly *logically* related, but does that mean that they are cognitively penetrated as well?

I think that the best way to set boundaries on this definition so that it does not result in absurd inclusiveness is to appeal to Fred Dretske's distinction between the causal and

---

know how to do things, or that perceptual content is conceptual. I think that many philosophers believe that the perceptual system must itself be cognitive in order to avoid the myth of the given or some other equally unhappy vestige of foundational epistemologies.

explanatory role of representational information. I hinted in this direction in my motor system example, but I would like to present a detailed account here.

Though Dretske's concerns are not with cognitive penetration, but rather with defining the limits of minimally rational behavior, it will be quite unproblematic to bring his theory to bear on questions of cognitive penetration. Dretske develops a principled criterion by which to determine when intentional states are related to behavioral output in such a way as to constitute an instance of minimally intelligent action. I will apply his standards for behavior to functional processing and, in this way, I will be able to ascertain when functional systems are cognitively penetrable. Dretske's standard for determining which connections to representational states are sufficient for minimally rational behavior will be my standard for determining which connections to cognitive states are sufficient for cognitive penetrability. Though Dretske refers exclusively to intentional or representational content and behavior, for my purposes I can expand this notion to include any cognitive content and its connection to functional processing. In this way, I can establish which functional systems are related to cognitive states such that they exhibit logical or semantic coherence of the right kind.

What's nice about Dretske's account is that it presents a principled way to differentiate between behavior that can be *described* as rational, which includes the behavior of thermostats and plants, and behavior which *is* rational, like the behavior of people and (some) animals. Co-opted for my own purposes, Dretske's notion of minimal rationality will allow me to cash out "logical relation" in such a way as to rule out the stomach's response to dieting as an instance of cognitive penetration. It will also allow logical deduction and other obviously cognitive processing to qualify as cognitively penetrable. And lastly, it will leave questions about the cognitive penetrability of the early perceptual processing systems open to further investigation.

Before moving on to Dretske's theory, we should notice that *minimal* rationality is not necessarily full-fledged rationality. As Dretske acknowledges, there are some disputes as to whether his standard gets us all the way to intelligent action.[60] But even if we need to add something to minimal rationality in order to get full-blown intelligence, what we do get with Dretske's theory is a non-arbitrary way to rule out many behaviors that would otherwise remain undifferentiated. That is, it may be that Dretske should rule out more than he does, but despite this, he provides us with a basic ability to rule out countless relationships that will not turn out to constitute even minimally rational behavior.

According to Dretske, the way to differentiate between minimally rational and thus, potentially intelligent behavior, and behavior of all other types, is to investigate what the role of intentional content is in determining and explaining behavior. For Dretske, we should be asking if an informational state is causally relevant or explanatorily significant to behavioral output.[61] If the representational content of an intentional state is explanatorily significant in accounting for the behavior of a system, then we have a case of minimal rationality.

For Dretske, representation is reducible to information, but not just any sort of information. If a state *S* has the function of indicating fact/state of affairs *f*, then that state is representational.[62] A representation is more than a carrier of information, since it not only carries information, but also represents it. The "function of indicating" does the heavy lifting here in order to differentiate, e.g., the fact that smoke indicates fire, and the fact that the beep of

---

[60] Fred Dretske, "Minimal Rationality," in *Rational Animals?* ed. Susan Hurley and Mathew Nudds, 107-8 (Oxford: Oxford University Press, 2006).

[61] "[I]f a structure's semantic character is unrelated to the job it does in shaping output, then this structure, thought it may *be* a representation, is not a belief. A satisfactory model of belief should reveal the way in which *what we believe* helps to determine *what we do"* Fred Dretske, *Explaining Behavior* (Cambridge, MA: MIT Press, 1988), 79.

[62] Dretske, *Explaining Behavior*, 51-64.

a smoke alarm indicates fire. Only the latter has the function of indicating fire, and so, only the latter qualifies as representational.

Possessing representational states, however, is not sufficient for being a minimally rational, intelligent or cognitive system. After all, many states have functions of indicating without having anything even remotely resembling rationality or intelligence.[63] Dretske's favorite examples of this sort of non-intelligent representational systems are thermostats and plants. Dretske points out that the internal states of such systems do in fact represent, they have indicator functions and this information does in fact have meaning (at least for us), but they do not having meaning that is related in the right way to the behavior of the system. That is, it is not the meaning of the state, but rather some other property of it, that is relevant to the character of the behavior of that system. For instance, the curvature of the metallic strip inside a thermostat does mean that the room is at a certain temperature, but it is not in virtue of this meaning that the thermostat kicks on the furnace. The physical bend of the strip and not the semantic properties that that bend represents cause the furnace to switch on. It is this very distinction between the meaningful content of a state providing reasons for a behavior, and some other property of that state being causally efficacious, that Dretske appeals to in order to differentiate between intelligent or minimally rational behavior and everything else.[64]

For Dretske, the content of a representational state cannot simply exist; it has to work. This distinction is clearly elucidated in Dretske's example of an opera singer's voice shattering a

---

[63] " Such tropistic behavior has a rather simple mechanical basis. And the blueprint for the processes underlying this behavior is genetically coded. The behavior is instinctive—i.e., not modifiable by learning. But it is not the simplicity of its explanation that disqualifies such behavior from being the behavior of interest in this study. *Reasons* are irrelevant to the explanation of this behavior, not because there is an underlying chemical and mechanical explanation of the movements in question (there is, presumably, some underlying chemical and mechanical explanation for the movements associated with all behavior), but because, although indicators are involved in the production of this movement, *what they indicate*—the fact that they indicate thus and so—is (and was) irrelevant to what movements they produce" Dretske, *Explaining Behavior*, 94.

[64] "It should be clear from my argument and examples that minimally rational behavior, is behavior that is explained by the meaning of content of the internal states that produce it" Dretske, "Minimal Rationality," 114.

glass.[65] The words that the opera singer sings have meaning, but that meaning is incidental to the glass's shattering. In contrast, when a rat moves in the direction of food, it's representational state "that the cheese is over there" explains why he moves to the left side of the maze. The content or meaning of the rat's representational state acts as a reason explaining his behavior, while the semantic properties of the opera singer's words are quite obviously irrelevant.[66]

By applying this criterion to cognitive penetrability, we are able to isolate which systems are related to cognitive states in a semantically coherent or logically relevant way. By ascertaining whether the meaning or content of a cognitive state is explanatorily relevant to the operations of a functional system, we can determine if it is cognitively penetrated or not. That is, if the processing of a functional system is not sensitive to the semantic content of cognitive states, then we can conclude that that system is not cognitively penetrable.

I think that now we are in a position to distinguish between the potential affects of cognition on perception, and the affects of cognition on one's stomach. Using Dretske's distinctions, we are in a position to say that whereas dieting, an intentional action that is logically related to all sorts of intentional states, does have the affect of shrinking one's stomach, it is not the case that the stomach is cognitively penetrated. This is because, while the semantic content of those intentional states play an explanatory role in the behavior of dieting, that same content does not play an explanatory role in the shrinking of the stomach.

Even though we have intentional states, which have semantic content that are responsible for the agent's actions, those states are not related to the changes in stomach size in an explanatorily relevant way. Though an agent certainly thinks "I want to be healthy," and this

---

[65] Dretske, *Explaining Behavior*, 179.

[66] As Daniel Dennett, *Content and Consciousness* (London: Routledge and Kegan Paul, 1969), 33 has noted, it is not only possible to attribute intentional states to animals, but it is impossible to explain animal behavior in the absence of such attributions. Such, has been the failure of behaviorism.

thought is responsible for the agent, say, cutting down portion size, it is not the semantic content, but rather some other chemical process, that explains why the stomach shrinks. Were the intentional state different or absent, one could still mimic the chemical environment resulting from the intentional choice to diet, and still get the affect of stomach shrinking. That is, the changes in the stomach can be explained in terms that have nothing to do with the reasons or intentional content of the desire to be healthy.[67]

So, I think the question to ask when it comes to cognitive penetrabiltiy is, "if the meaning or cognitive content of a state were different, could the response by a system be the same?" If the answer is "no," then we have an instance of cognitive penetration. Such a test will allow us to determine when functional processes are logically related to cognitive content in explanatory significant ways, and thus to determine when changes internal to a functional system are the result of cognitive penetration.

In this way, we can differentiate between mechanistic, automatic, tropistic changes and the results of cognitive penetration. By applying Dretske's principle, we will be able to determine if the changes in processing that occur at a time, or over time, are the result of cognition, properly so called. That is, by appealing to Dretske's distinction, we can cash out semantic coherence or logical relatedness as a relation to states whose content, not just existence or instantiation, plays an explanatory role in changes, which occur in a system.

4. The flexibility of automatic systems:

4.1 Automatic systems and cognitive penetration:

Finally, let us turn to Fodor's third purported reason for denying that long-term cognitive penetration is truly an instance of cognitive penetration—because early perceptual systems are

[67] This does *not* entail that cognitive states are not physically instantiated. It simply means that the description of the physical properties of a cognitive state is not sufficient for providing an explanation of behavior.

automatic, any long-term cognitive penetration is not really cognitive. In this section, I hope to show that this objection is based in confusion.

When we think of automatic systems, we think of rigid, inflexible, programmed, reflex-like systems that are triggered in the absence of choice or volition. It would seem to follow that if early perceptual systems are automatic, then they must be immune to the impacts of cognition. This follows because cognitive penetrability requires a sensitivity and a capacity for change that automatic systems, as described above, do not possess. In this section, I will argue that it is possible for a system to be both automatic and cognitively penetrable. If we are careful to distinguish between the level of information processing and the level of rule implementation, then it will be clear that automaticity and flexibility are not mutually exclusive. I will argue that information may be processed automatically at a given time, while the rules which determine the processing of that information may be responsive to cognitive guidance over time. To be automatic at a time does not entail diachronic inflexibility.

A thermostat set at 68 degrees will function differently from that same thermostat when it is set at 70 degrees. Both functions, however, are equally automatic. While the first function can be described as implementing the rule "kick on furnace when room temperature falls below 68 degrees," the other should be described as, "kick on thermostat when room temperature falls below 70 degrees." The rule that the thermostat follows has changed though the automatic processing of the thermostat has remained the same. Certainly, this is an extrinsic type of flexibility, but it does elucidate the fact that automatic processing does not entail that a function must remain identical at all times in order to remain automatic.

One may argue, however, that the function of the thermostat has not changed—only its threshold has. One may legitimately claim that the rule implemented in both cases is, "kick on

furnace when room temperature falls below the temperature set" and thus, it is only the setting and not the function that undergoes any type of alteration. One may claim that the automaticity of the thermostat does not allow for a change in function.

So, let's look at a different example. Let's look at an updated computer program. Since the updating of a computer program consists of replacing certain sections of a program with updated instructions, we can rest assured that what we have here *is* a change in function. This should be clear since the next time an input ends up at the updated section of the program, it will be processed according to a new set of rules. But this change in function does not make the program itself any less automatic. The rules according to which the input will be processed may have changed, but the input will still be processed automatically. It should be clear that a change on the level of rules or procedures can occur without spiraling into unbounded flexibility.

At this point, it may be objected that I have not described a change in a program but an entirely new program altogether. That is, it is possible that I will run into a problem of individuation and identity over time here. After all, one may claim that any change in the procedural processing of a program entails that the identity of the program perishes. One may claim that the program ceases to be the same program once the rules by which it functions have changed. As a result, the above example would fail to display an instance of a change in function with the retention of automaticity.

My response to this concern is rather unexciting. I suggest appealing to a pragmatic criterion for individuating programs and functions and really, just about everything else. Persons, objects, thoughts, processes, events, etc., should all be individuated according to a pragmatic principle. I have found no arguments in discussions of personal identity or in debates about the identity of objects over time to convince me otherwise. As such, this objection falls

short. Or rather, it overshoots. Nothing is safe from the problem of numeric identity through transformation and, as such, that this particular issue is not immune to it either, should not give us much cause to fret.

Furthermore, the possibility of restructuring the computational procedures inherent to a cognitive system is not just a theoretical possibility, but a hot bed of empirical research.[68] A promising line of thought posits that automatic behavior is the result of changes in the strategy of processing and not just in the speed of processing. According to this theory, expert behavior differs from the laypersons not just in how quickly an expert can process information, but in the actual algorithms that the expert implements in order to process that information. It is hypothesized that cognitively sensitive changes are central to the implementation of automatic behavior. Of course, this is not an argument for the cognitive penetrability of automatic early perceptual systems, but it does encourage us to rethink our pre-theoretical notions about the relationship between automaticity and cognition.

And even further, there is nothing in principle to prevent conceptual or cognitive tasks from being automatically instantiated. That is, there is no reason why automatic should mean noncognitive, in the first place. For example, the search performed by an expert chess player or the arithmetic calculations of a mathematician can both function in a perfectly automatic fashion. As Jesse Prinz has argued,

---

[68] L. Saling and J.G. Philips, "Automatic Behavior: Efficient Not Mindless," *Brain Research Bulletin* 73 (2007): 1–20; P.W. Cheng, "Restructuring Versus Automaticity: Alternate Accounts of Skill Acquisition," *Psychological Review* 92 (1985): 414-423; J.J. Staszewski, "Skilled Memory and Expert Mental Calculation," in *The Nature of Expertise*, ed. M.T.H. Chi, R. Glaser and M.J. Farr, 71-128 (Hillsdale, NJ: Erlbaum, 1988); S.B. Smith, *The Great Mental Calculators: The Psychology, Methods, and Lives of Calculating Prodigies, Past and Present* (New York: Columbia University Press, 1983); R.W. Doerfler, *Dead Reckoning: Calculating without Instruments* (Houston: Gulf Publishing, 1993); K.Anders Ericsson and Neil Charness, "Expert Performance: Its Structure and Acquisition," *American Psychologist* 49 (1994): 725-74; W.G. Chase, and H.A.Simon, "The Mind's Eye in Chess," in *Visual Information Processing*, ed. W. G. Chase, 215-81 (New York: Academic Press, 1973); D. LaBerge and S.J. Samuels, "Toward a Theory of Automatic Information Processing in Reading," *Cognitive Psychology* 6 (1974): 293-323; G.D. Logan, "Automaticity, Resources, and Memory: Theoretical Controversies and Practical Implications," *Human Factors* 30 (1988): 583–598.

> Semantic priming is mandatory, but it taps into conceptual knowledge…everyone agrees that some mental processes are automatic, but most mental capacities seem to integrate automatic processes with processes that are controlled. For example, we form syntactic trees automatically, but sentence production can be controlled by deliberation. Likewise, we see color automatically, but we can visually imagine colors at will. The automatic/controlled distinction cannot be used to distinguish systems in an interesting way. [69]

As such, the automatic/cognitive distinction is hardly instructive in distinguishing between intelligent and tropistic tasks or functions. The idea, then, that automaticity rules out cognitive penetration is untenable.

To conclude, it ought to be clear that there are many reasons to doubt that the automaticity of a system entails its exclusion from the category of the cognitive. A system may automatically process conceptual information, or it may possess relatively flexible and responsive structural rules and procedures whose flexibility is not inherited by the inputs that are processed by them. Both of these considerations are sufficient to rule out the presupposition that automatic processing is necessarily unintelligent.

4.2 Flexibility, generally

To conclude this chapter, I'd like to make some general remarks about the flexibility of cognitive systems. The purpose of these remarks is not to lay the groundwork for an argument concluding that perceptual systems are indeed cognitive, but rather to gesture at the possibility that our pre-theoretical assumptions concerning cognition are unrealistic. The upshot of these considerations will relate less to perceptual systems and more to the determination of which states and actions we are willing to classify as cognitive. This will prove important in chapter five.

Fodor has stated,

---

[69] Jesse Prinz, "Is the Mind Really Modular?" in *Contemporary Debates in Cognitive Science*, ed. Robert J. Stainton, 25 (Malden, MA: Blackwell, 2006).

> We have only the narrowest of options about how the objects of perception shall be represented, but we have all the leeway in the world as to how we shall represent the objects of *thought;* outside perception, the way that one deploys one's cognitive resources, is, in general, rationally subservient to ones utilities.[70]

It is clear from this statement that Fodor would like to distinguish thought, in all its freedom, from perception, that prisoner of transduction. And Fodor is not alone, for it is often considered a truism that thought—conscious, propositional, linguistic, adult thought, exhibits a type of flexibility that other systems do not. The idea is that you can think however you want, but you cannot digest or perceive however you want.

Now, I'm quite sure that Fodor overstates the case and I believe that Moore's Paradox is quite instructive in pointing out how. Famously, G.E. Moore argues that it is impossible to say that one believes that *p* and also to say that one believes that *p* is false.[71] That is, to believe that *p*, commits one to the belief that *p* is true. For my purposes, limitations on the freedom and flexibility of thought reveal themselves just as we begin asking persons to believe things that they believe are false, or vice versa, to stop believing things that they believe are true. Here's an exercise: Believe that New York is the capital of France; believe that chocolate cake is health food, or that Obama is not a better speaker than Bush. The thing is that you can't do it. Why? Because your beliefs, it turns out, are not simply the products of your will![72]

Though it is clear that one's beliefs are not the direct result of one's will (as any person that has ever attempted to change the way she thinks can tell you), it may be claimed that the entertaining of thoughts is all that we need in order to ensure the ultimate freedom of the cognitive realm. It may be claimed that what we choose to think about is our choice, even if what we end up believing is not. Though it is notoriously difficult to control what one thinks

---

[70] Fodor, *The Modularity of Mind,* 55.
[71] George Edward Moore, "A Reply to My Critics," in *The Philosophy of G. E. Moore* ed. P. A. Schilpp, 535-677 (Evanston, IL: Northwestern University Press, 1942).
[72] Just think about how many basic beliefs we possess: beliefs about gravity, solid objects, movement, etc.

about (think meditation), for the sake of argument, let's admit that this is no great feat. What remains then is not *what* one chooses to entertain, but rather *how* one chooses to entertain it.

However, it turns out that the *how* fares far worse then the *what*. One's line of reasoning, or one's process of thinking is almost entirely opaque. Like the processes of the early input systems, most of us have no idea how we go from one thought to another. I cannot tell you which processes I undertake to transform my thought about Cinderella's pumpkin into a thought about pumpkin pie as opposed to a thought about Clydesdale horses. As Ryle points out:

> There are many classes of performances in which intelligence is displayed, but the rules or criteria of which are unformulated. The wit, when challenged to cite the maxims, or canons, by which he constructs and appreciates jokes, is unable to answer. He knows how to make good jokes and how to detect bad ones, but he cannot tell us or himself any recipes for them… Some intelligent performances are not controlled by any anterior acknowledgement of the principles applied in them.[73]

And because we do not know the procedures by which we process our thoughts, we are unlikely to find any formula by which to manipulate them. We are in a sort of double bind. We don't know how we get from one thought to another and we don't know what ought to be done in order to change or control the process.

The take-away point here is simple: there are no fully flexible cognitive systems (not any that we know of anyway). As such, not possessing the property of infinite flexibility cannot entail being excluded from the category of the cognitive.[74] If this is our standard then nothing qualifies as cognitive. But surely this is too high a price to pay for unexamined ideals.

5. Conclusion:

[73] Gilbert Ryle, *The Concept of Mind* (Chicago: The University of Chicago Press, 1949), 30.
[74] David Papineau and Cecelia Heyes, "Rational or Associative? Imitation in Japanese Quail," in *Rational Animals?* ed. Susan Hurley and Mathew Nudds, 194 (Oxford: Oxford University Press, 2006) make a similar point when considering which states qualify as representational.

To conclude, by elucidating the commitments that follow from the definition of cognitive penetration, I have argued that Jerry Fodor's position concerning the status of long-term cognitive penetration is hardly decisive. I have argued that since cognitive systems are not identical to cognitively penetrable systems, an objection to the authenticity of long-term cognitive penetration cannot rely on the fact that the internal connections, which result from long-term cognitive penetration are not themselves cognitive. The fact is that in order to constitute an instance of cognitive penetration it is not required that such connections are cognitive, just so long as they are related to cognitive states in the right way.

I have also argued that if cognitive states can be nonlinguistic and nonpropositional, then the fact that the states that impact perceptual processing in the long term are neither linguistic nor propositional will not entail that they are not cognitive, and thus, they cannot be ruled out as potentially responsible for cognitive penetration. So, if Fodor was arguing that diachronic cognitive penetration is not really cognitive penetration because the states that affect perceptual processing are not propositional, then he must also prove that all nonpropositional states are not cognitive. It is with this very issue that I will endeavor to respond to Fodor.

I have also attempted to cash out the notion of a semantically coherent or logical relationship between cognitive states and a functional system. I did this by relying on Dreske's theory of minimal rationality. As I develop my theory, this notion will prove critical in order to discern when nonpropositional cognitive states are related to early perceptual systems such that they manifest instances of cognitive penetration.

I ended this chapter by arguing that automatic processing is not necessarily inflexible. I have argued that the level of rules and procedures can be flexible in such a way that does not compromise the automaticty of information processing internal to a functional system. This

argument was a response to the proposed objection that long-term cognitive penetration is not really cognitive penetration, because early perceptual systems are automatic.

CHAPTER 3

1. The Problem

In the previous chapter of this dissertation, I presented cursory evidence supporting the position that a Language of Thought-based conception of cognition is not comprehensive. I indicated that cognition is likely not exhausted by the category of intentional, discursive, propositional, truth-functional states. Here, I would like to expand on this idea in order to build a foundation for a positive account of diachronic cognitive penetration.

In the following chapters, I will argue that if early perceptual processing is impacted in the right way by intelligent abilities or skills, then this interaction constitutes a genuine instance of diachronic cognitive penetration. Such a position does not undermine Jerry Fodor and Zenon Pylyshyn's claims that propositional states do not penetrate the functioning of early perceptual systems, but it does open up the possibility for a new approach to cognitive penetration. Though I will not argue that intentional states impact the functioning of early perceptual systems, I will promote the possibility that cognitive behavior such as the instantiation of skills, may. This means that if perceptual systems are supposed to be modular, and if modular means that no species of cognition can impact those systems, then the possibility of skill caused diachronic cognitive penetration will undermine the modularity of the early perceptual systems. As I stated earlier, such cognitive penetration is not a free-for-all. The affects of skill on early perception will amount to a limited and systematic relationship between intelligent action and the processing of qualitative properties.

So, though I agree with Fodor that intentional states are unlikely to penetrate early perceptual processing, I part ways with him in denying that propositional, conceptual, intentional

thought exhausts the realm of the cognitive. The application of intelligent abilities, I argue, is an instance of nonpropositional, practical, embodied intelligence. It is this type of cognition that I believe may, in fact, affect early perceptual processing. It is in virtue of the distinction between propositional thought and intelligent ability that I can do justice to the fact that some, but not all, cognitive states or events are capable of impacting perceptual processing. That is, by distinguishing embodied, learned abilities from propositional thought, I can explain why these types of intelligent behaviors may have an impact on the qualitative character of a perceptual event while holding that propositional knowledge does not.

I submit that as a result of the instantiation of skills, changes in perceptual processing may occur. Or more precisely, my point is this: if intelligent abilities hold the right kind of logical relationship to perceptual processing, then this will count as a legitimate instance of cognitive penetration.

Crucially, the cognition involved in skill is not reducible to the intelligence of propositional states that are often associated with skillful performance. Though skills are intimately connected to beliefs, desires and knowledge, it is not the case that the intelligence of skill is reducible to the intelligence of the intentional states that are related to them. Following Gilbert Ryle, I will argue that though propositional states will certainly count as cognitive, the application of intelligence in action manifests a distinct, but equally legitimate cognitive category.

The purpose of this chapter will be to argue that the successful exercise of skills is not reducible to the intelligence of knowing theories or facts about them. In arguing that skill is a distinct type of cognitive event, I will make headway towards establishing the possibility that practical intelligence may be able to cognitively penetrate perception while maintaining that

propositional thought cannot. By distinguishing *knowing-how* from *knowing-that*, I will be in a position to build a positive theory of skill and practical intelligence in chapters four and five. Here, I will be concerned to say what skill is *not*. In the final chapters of this dissertation, I will say what it *is*.

It is the purpose of this chapter to provide a decisive argument for non-propositional intelligence. I argue that a complete account of propositional knowledge cannot explain the intelligent application of propositional thought in action. I do this by first presenting Gilbert Ryle's argument in favor of a distinction between knowing-that and knowing-how. Next, I rehearse and respond to Jason Stanley and Timothy Williamson's objection to Ryle. In order to clarify the disparity between propositional knowledge and intelligent ability, I present an example where this distinction is clearly illustrated. From this example, three principles follow: (1) propositional knowledge is not sufficient for know-how, (2) differences in know-how are not proportional to differences in propositional knowledge, and (3) knowing-how to *a* requires the ability to *a*, but it does not require the ability to *a* now. I develop principles (1) and (2) here[75] and conclude by clarifying what know-how is not. In this final section, I explicitly argue that know-how cannot be an instance of practical reasoning, it is not conceptual or propositional, and it is not directed at truth.

2. Ryle on knowing how

In *The Concept of Mind*, Gilbert Ryle argues that there are two types of knowledge: knowledge-that and knowledge-how. Knowledge-that is the type of state that everyone from Plato to Fodor takes to be real knowledge. It is propositional and when it hits its mark, it is true and justified. Distinctively, Ryle argues that if anyone is to know *that* anything is the case, then another kind of knowledge must be operative. This other type of knowledge is practical

[75] I address principle (3) in chapter four.

knowledge. It is more than just knowing the fact that *p* is true—it is knowing how to appropriately apply the criteria by which to determine whether *p* is true; it is this kind of knowledge that allows one to apply the knowledge that *p* is true into practice. Ryle states that, "to be intelligent is not merely to satisfy criteria, but to apply them."[76] That is, intelligence is not simply a matter of validly deducing what to do and when, but the capacity to execute that judgment in action.

Ryle's intent in this chapter of *The Concept Of Mind* is to destroy the intellectualist legend, which holds that the only things that count as intelligent or mental, are those things that contribute to the knowledge of truths. Ryle is adamant about re-introducing behavior, ability, activity and skill into the realm of the mental, and hence for him, the realm of the cognitive. He states,

> [T]here are many activities which directly display qualities of mind, yet are neither themselves intellectual operations nor yet effects of intellectual operations. Intelligent practice is not the step-child of theory. On the contrary, theorizing is one practice amongst others and is itself intelligently or stupidly conducted.[77]

By shifting focus from contemplative states of mind to the intelligent actions of embodied persons, Ryle is attempting to introduce a role for the mental that is both natural and observable.

For Ryle, know-how is intelligent ability and abilities are dispositions. Such dispositions are not innate, but rather learned through practice and characterized by intelligence. Ryle takes the following to be paradigmatic instances of knowing-how: playing chess, appreciating jokes, performing surgery, speaking grammatically and reasoning well. None of these, he claims, are reducible to knowing the rules governing the successful instantiation of these activities. Ryle

[76] Gilbert Ryle, *The Concept of Mind* (Chicago: The University of Chicago Press, 1949), 28.
[77] Ryle, 26.

states, "ability at reasoning is exhibited in the construction of valid arguments and the detection of fallacies, not in the avowal of logician's formulae."[78] After all, he argues, one can know all the rules, and still not be able to perform the requisite activity successfully. For Ryle, it is not the case that one has a theory of reasoning and thus reasons well by applying that theory, but rather, the reasoning itself—the manifestation of the ability to reason—is a cognitive operation irreducible to any prior set of propositions. In short, theory does not birth practice.

Ryle argues that if the application of intellect in practice were reducible to propositional knowledge, then a regress would swiftly ensue.[79] Ryle argues that it is impossible that knowing how to do something requires first thinking of the rule, which governs the behavior of how to do it. For, if knowing-how required contemplating a proposition in order to apply it, then one would also need to contemplate another proposition in order to know how to apply the proper contemplating of the first proposition, and so on *ad infinitum*. Ryle states the regress like this:

> The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle.[80]

This regress strikes a decisive blow to the intellectualist legend. It mutes the notion that intelligent behavior can be reduced to the justified believing of truths. The regress proves that it is impossible to reduce reasoning well, or playing chess, or telling funny jokes, exclusively, to the possession propositional knowledge.

---

[78] ibid, 49.

[79] Aristotle makes a similar point in *The Nicomachean Ethics* (Oxford: Oxford University Press, 1988), Book III 1113 a1 when he argues that the mean of virtue is the result of perception, not reasoning: "…nor indeed can the particular facts be a subject of [deliberation],e.g. whether this is bread or has it been baked as it should; for these are matters of perception. If we are to be always deliberating, we shall have to go on to infinity." We should note that Aristotle is not making a claim about raw sensations being necessary for a foundationalist theory of epistemology. Whether bread is baked as it should be is quite obviously not a matter of raw percept.

[80] ibid, 30.

3. Ryle: criticism and response

For years, the distinction between knowing-that and knowing-how was accepted and appealed to, but neither questioned nor developed. Recently, however, a debate surrounding Ryle's distinction has emerged. Jason Stanley and Timothy Williamson's 2001 article "Knowing How" has been uniquely responsible for spearheading this resurgence. In that article, Stanley and Williamson attempt to both undermine Ryle's regress argument and provide an alternate theory of knowing-how.[81] Fortunately for Ryle, they neither present a decisive blow against his regress argument, nor provide a satisfactory explanation of intelligent abilities as a species of propositional knowledge.

Stanley and Williamson claim that Ryle's regress does not work because in order for it to begin, each knowing-how must be an instance of an intentional action. But often the contemplating of propositions is not intentional at all. And when the contemplation of a proposition is not an intentional event, it is not the type of thing that can fuel Ryle's regress. Stanley and Williamson's argument goes like this: First, they attribute the following two premises to Ryle.

(1) If one F's, one employs knowledge how to F
(2) If one employs knowledge that p, one contemplates the proposition that p.

> If knowledge-how is a species of knowledge-that, the content of knowledge how to F is, for some φ, the proposition that φ(F). So the assumption for reductio is:

RA: knowledge how to F, is knowledge that φ(F)

> Furthermore, let 'C(p)' denote the act of contemplating the proposition that p. Suppose that Hannah F's. By premise (1) Hanna employs knowledge how to F. By RA, Hanna employs knowledge that φ(F). So, by premise (2), Hannah C(φ(F))s. Since C (φ (F)) is an act, we can reapply premise 1, to obtain the conclusion that

[81] See Jason Stanley and Timothy Williamson, "Knowing How," *Journal of Philosophy* 98 (August, 2001): 411-44.

> Hannah knows how to C(φ(F)). By RA, it then follows that Hannah employs that knowledge that φ(C(φ(F))). By premise (2) it follows that Hanna C(φ(C(φ(F))))s and so on.[82]

Now, Stanley and Williamson argue that this regress doesn't get off the ground because, as they rightly point out, not all doings are instances of knowing-how to do. Their example of Hannah digesting food, but not knowing-how to digest food, is a good example of a nonintentional action.[83] They correctly stipulate that not any action or ability will start a regress, but only *intentional* actions or abilities. From here, Stanley and Williamson move to consider whether the contemplation of a proposition is always an intentional act. In order to do this, they appeal to Carl Ginet, who argues that not all instances of knowing-that are accompanied by the contemplation of a propositional thought. He states:

> I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition.[84]

Ginet's claim, if true, would arguably render premise (2) false. But, rather than insist that premise (2) is false, Stanley and Williamson choose to argue that contemplating a proposition, in some cases, is not an intentional action at all. That is, contemplating the proposition that *p* is more like Hannah's digesting food than Bill's searching for mistakes in his article. Thus, they render such actions improper substitutions for F in premise (1). The force of the regress is deflated.

---

[82] Stanley and Williamson, 414.
[83] Alva Noë, "Against Intellectualism," *Analysis* 65 (2005): 279 argues that to say that digesting food is something Hannah does, is suspect in the first place. It is not an act at all, but a happening. Hannah does not DO anything, well, fast or otherwise.
[84] Carl Ginet, *Knowledge, Perception and Memory* (The Netherlands: D. Reidel Publishing Company, 1975), 7.

There is a lot that is funny about this argument, but I will limit myself to the most immediate problems. First, Ginet's observations about not needing to contemplate a proposition in order to manifest propositional knowledge cannot lead us to conclude anything about the metaphysical nature of contemplating propositions beyond the fact that phenomenologically it is does not always *seem* to us that we are contemplating a proposition when we are manifesting propositional knowledge. As Alva Noë points out, the most we get out of Ginet's claim is that there is not always *conscious* contemplation of a proposition whenever knowledge that *p* is manifest.[85] But this phenomenological absence is hardly a decisive position from which to describe the nature of propositional knowledge. It certainly does not entail that nonconsious contemplation is not intentional.

Quite simply, Stanley and Williamson conflate consciousness with intentionality. They are not entitled to the move from "it does not seem to me that I am contemplating the proposition that *p*" to "the contemplation of the proposition that *p* is not an intentional action." And without this move, Ryle's regress stands unharmed.

Additionally, to say that knowing-how is non-intentional rather than non-propositional proves to have potentially devastating results for Stanley and Williamson's positive account of abilities. Stanley and Williamson's overall goal is to show that knowing-how claims can be cashed out in the language of knowing-that. However, the kind of phenomena that could potentially be explained in terms of knowing-that would need to be intentional. Unambiguously, if the kinds of actions that knowing-how is meant to explain could really be cashed out in terms of knowing-that, then those phenomena would be intentional actions and, hence, Ryle's regress would begin. On their own revised account of know-how, this is exactly what happens.

[85] " Ryle can accommodate Ginet's observation by countenancing the possibility that not every act of contemplating a proposition is performed consciously. To say that it is or could be performed unconsciously is not to say that it is not the sort of thing that could be performed intentionally" (Noë, "Against Intellectualism," 282).

So, the only thing that Stanley and Williamson have successfully argued is that contemplating, reasoning, inferring etc., are not always intentional actions. They are more like digestion—things that happen to one but not things that one does. But this clearly won't do because these are the very phenomena that Stanley and Williamson seek to explain in terms of propositional knowledge. They claim that knowing-how to ϕ is really an instance of knowing that *w* is a way to ϕ under a practical mode of presentation. But knowing that *w* is a way to ϕ under a practical mode of presentation seems squarely intentional and, thus, the type of thing that fuels Ryle's regress.[86] As such, Stanley and Williamson's critique of Ryle fails.

Further, Stanley and Williamson's positive account is not only guilty of starting the very regress that they claim knowing-how, construed as propositional knowledge, does not start, but it also fails to explain the exact phenomena that knowing-how is supposed to describe: the application of knowledge in practice. Stanley and Williamson propose that the proper construal of knowing-how is as follows:

> Hannah knows [PRO 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.[87]

This account is thoroughly unsatisfactory since it appeals to the very notion that it is trying to deny. For Ryle, the distinction between knowing-that and knowing-how is the distinction between thought and ability. Ryle's point is that abilities cannot be reduced to the

---

[86] "Suppose that s φ's. Then by (1), she employs knowledge how to φ. So by RA, there is s way w such that she knows that w is a way to φ. By (2), she contemplated the proposition [*w* is a way to φ]. By (1) again, she employs knowledge how to contemplate the proposition [*w* is a way to φ]. So by RA, there is another way *w** such that she knows that [*w** is a way to contemplate the proposition [*w* is a way to φ]. And *ad infinitum*...in plain English, this says that know-how *does not always* consist of propositional knowledge" (John N.Williams, "Propositional Knowledge and Know-How," *Synthese* 165 (2008): 112). See also John Koethe, "Comments and Criticism: Stanley and Williamson on Knowing How," *Journal of Philosophy* 99 (2002): 325-328.

[87] Stanley and Williamson, 430.

possession of a certain set of facts, since the possession of knowledge does not itself dictate its intelligent or stupid application. What Stanley and Williamson have done is to define abilities as knowing facts under a *practical mode of presentation.* But that practical mode of presentation is exactly what makes knowing-how different from knowing-that. It is this practical mode—this manifestation of knowledge in action that itself expresses various degrees of intelligence, which cannot be reduced to knowledge of a particular set of propositions.

Thus, it remains to be seen *how* propositional knowledge is –or, should be understood to be – instantiated to various degrees of success and elegance under a practical mode of presentation. Stanley and Williamson attempt to account for this critical issue by saying that propositional knowledge is put into practice by being contemplated under a practical mode of presentation,[88] but any account of entertaining a proposition under a practical mode of presentation is conspicuously missing from their article. The closest they get is to say that such an account is surely possible. As Noë writes:

> [I]t is difficult to see how the positive analysis offered by Stanley and Williamson entails the falsehood of Ryle's distinction between knowing how and knowing that… Stanley and Williamson's preferred account doesn't eliminate the distinction, or give anyone committed to it a reason to give it up; it merely relocates it.[89]

The point is not that propositional knowledge is not often or cannot ever be a factor in the intelligent exercise of abilities. It is fairly uncontroversial that many abilities require at least some sort of basic propositional knowledge for their successful performance. Rather, the point is that if a complete account of propositional knowledge is given, it will not constitute an exhaustive account of intelligent abilities. Applying knowledge appropriately is not the same as possessing it.

---

[88] See Williams, "Propositional Knowledge and Know-How" and Koethe, "Comments and Criticism: Stanley and Williamson on Knowing How" for similar points.
[89] Noë, "Against Intellectualism," 287.

This does not, of course, mean that one cannot give an explanation of intelligent ability and categorize it under the heading "propositional knowledge". One could place "knowing-how" under the heading "knowing-that," with the qualification that it is a knowing-that of a very special kind. But then one would still be left with the responsibility of presenting a satisfactory account of what makes this type of propositional knowledge so special. The fact remains that disputes about nomenclature hardly address the substantive issue here. The issue is that knowing the rules governing an action does not entail an ability to successfully perform that action, even if one has the physical strength, mental prowess, or opportunity to do so. This means that whatever category we want to place knowing-how into, knowing-how will still require explanations that are distinct from those accounting for the possession of propositional knowledge.

4. Bela and Mary Lou: an example

Bela Karoli is the world-famous coach of several gold medal-winning women gymnasts. Bela Karoli knows the rules governing the skills that he teaches his athletes, but he is unable to perform these skills himself. It is not the case that Bela could once perform these skills, but is too fat or too old to perform them now. He never knew how to perform those skills and neither a good diet nor a good time machine could change that fact.[90] Bela Karoli knows the rules governing the skillful performance of, e.g., a standing layout on beam, and he also knows how to express that knowledge in such a way that his gymnasts can apply it to their own learning and performance. What he himself does not know, however, is how to perform a standing layout on beam.

[90] I mention this because usually examples of coaches possessing abilities are given in terms of them once having had the ability, but no longer being able to successfully perform it. But the Bela is distinct, since he never had the ability in the first place.

Importantly, the missing link is not just that Bela doesn't know what it's like to perform this skill. We can imagine that Bela is put into a gymnastics-simulator and so has the experience of feeling what it's like to perform the skills he teaches.[91] Upon, exiting the simulator, Bela has all the propositional knowledge he can ever have about standing layouts on beam, he has felt what it is like to perform the skill, and for good measure, let's make him young and fit, too. Still, none of this will amount to him having the ability to perform the skill. There is something more to knowing how to do a standing layout on beam, and that something else, is what Bela lacks.

Imagine that we ask Bela the following question, "Do you know how to do a standing layout on beam?" An acceptable response would be, "I know how to teach someone how to do a standing layout on beam, but I do not know how to do one myself." Compare Bela's answer to Mary Lou Reton's. If we asked her, "Do you know how to do a standing layout on beam?" she should answer, "I know how to do a standing layout on beam, but I cannot do one now, since I am too old and out of shape." Bela cannot perform the skill because he has never learned how to perform the skill. Mary Lou, on the other hand, cannot instantiate her ability due to the fact that the conditions of her physical body act as an obstacle to her performing it.[92]

For Mary Lou, the inability to perform a standing layout on beam is a matter of opportunity; it is the inability to put her ability into practice now.[93] But for Bela, it is a matter of lacking the ability, period. He cannot put the ability into practice not because he lacks the

---

[91] I have in mind here a simulator where Bela is passive and he gets to experience performing the skill, without any control or feedback affecting the experience. We could of course imagine a simulator where Bela is controlling his movements and as a result, he may learn how to perform the skill. That is, after all the point of many simulators (flight and otherwise), but that is more than just getting to feel "what it's like."

[92] See Ginet, 1975, Stanley and Williamson, 2001 and David Carr, "Knowledge in Practice," *American Philosophical Quarterly* 8 (January, 1981): 53-61 have claimed that know-how is not ability because there are cases where people know how to do *a*, but don't have the ability to do *a*. This is an important point, but it is even more important that we are clear on what it means to say that S does not have the ability to do *a*.

[93] We can imagine that Mary Lou has a broken ankle, if that will help make this point clearer.

opportunity, but because he lacks the ability. It is vital for us to differentiate between the ability that Bela doesn't have and the ability that Mary Lou lacks.[94] This difference helps us elucidate the necessary relationship between knowing-how and ability.[95]

Importantly, the difference between Bela and Mary Lou will not be obliterated if we give Bela any more propositional knowledge or a more fit body. And crucially, it is not essential or even likely that Mary Lou possesses anywhere near the amount of propositional knowledge that Bela does, and yet, she still knows how to do something that Bela does not. That is, Mary Lou, if young and fit, would be able to perform a standing layout on beam, despite the fact that she has less propositional knowledge than Bela and the case would not change, even if Bela was as young and fit as Mary Lou.

We should glean at least three principles from the above example: (1) Propositional knowledge is usually not sufficient for ability, (2) differences in expertise are not always proportionate to differences in propositional knowledge, and (3) knowing-how to *a* requires the ability to *a*, but it does not require the ability to *a* now. I will address principles (1) and (2) here and principle (3) in chapter four, section two.

4.1 Propositional knowledge is hardly ever sufficient for ability

The above example should make it abundantly clear that propositional knowledge is not sufficient for ability or know-how. After all, Bela knows all there is to know about the facts and rules governing the performance of a standing layout on beam, but he still does not know how to do a standing layout on beam. That is, he does not have the ability to perform the skill. The discrepancy between knowing the rules governing the performance of *a* and being able to

[94] I will discuss this difference in detail in association with criterion (2) of ability-possession.
[95] This is important because many intellectualist proponents tend to deny that there is a necessary connection between ability and know-how and this is how they justify treating Bela's knowledge as an instance of knowing-how. The distinction that I make here highlights why such a move is unwarranted.

successfully perform *a* is a fairly widely observed phenomenon inspiring the adage, "those who can't do, teach." The main point here is simply that possessing knowledge of rules and facts is not identical to possessing the ability to put those facts and rules into practice effectively. If it were, then anyone that had the requisite knowledge and opportunity could successfully instantiate the ability. But this is clearly not so, as the above example elucidates. Importantly, what is missing in these cases is not more facts or rules. And the reason for this is that the ability to apply propositional knowledge amounts to something more than being in possession of it.[96]

It may be argued that there are some instances where the possession of propositional knowledge *is* sufficient for ability. I have not encountered a convincing instance of such a case, but I think it will be useful to examine some proposed examples.

In "Propositional Knowledge and Know-How," John Williams presents the example of Sam who has been told the way to get to the post office. Sam has been given directions, and simply in virtue of being in possession of these new propositional facts, Williams claims that Sam now has the ability to get to the post office; he claims that these facts are sufficient for ability.[97] Now, I find this example a little hard to accept. It seems quite obvious that simply having the requisite propositional knowledge (that to get to the post office one must take a right, walk two blocks, take a left, and then an immediate right) does not guarantee anyone's ability to get to the post office.[98] After all, if it did, it would mean that anyone that had this propositional

[96] As an example of an abundance of rules without a corresponding ability, Ryle says the following " a boy is not said to know how to play [chess], if all he can do is to recite the rules accurately…a foreign scholar might not know how to speak grammatical English as well as an English child, for all that he had mastered the theory of English grammar" (Ryle, 41).

[97] Williams, 115.

[98] I am ignoring physical obstacles like not being able to walk and not having a wheel chair, or a street being closed because of a parade as the reasons one couldn't get to the post office. I am assuming that one could, if she knew how to follow directions, get to the post-office. I am questioning, however, whether physical opportunity plus propositional knowledge is sufficient for successfully following the directions.

knowledge could get there. But, that is simply not true. People with good directions get lost all the time.

What if Sam always confuses his right for his left, or he doesn't know which instruction to follow first and which second, or he doesn't know that turning right requires turning right on a street and not into an ally or a driveway? It seems that Sam has to know a lot about the world and how to act in it in order to follow directions. That is not to say that none of Sam's knowledge is propositional, but it is to say that it is far from obvious that *all* of it is. Williams' point is that sometimes just having propositional knowledge is enough to guarantee ability, but surely, getting to the post office requires a whole host of other basic capacities that are not obviously or easily reducible to propositional knowledge. For these reasons, his example fails.

John Bengson and Marc Moffet also claim that the possession of certain concepts entails particular abilities. They say that, e.g., having the concept of addition entails that one is able to add. They state,

> [T]he sort of reasonable mastery of these concepts presupposed by know-how attributions require that the subject possess the corresponding abilities…if our treatment of the puzzle is correct, then know-how depends on reasonable conceptual mastery, and thus a certain sort of understanding, in an important range of cases…reflection on the cognitive nature of intentional action and the fact that certain misunderstandings undermine know-how even when the corresponding abilities are in place show that this connection can be generalized, thereby allowing a general intellectualist analysis of know-how in terms of understanding. In short, the key to an adequate philosophical theory of know-how is not the relation between know-how and ability, but rather the connection between know-how and concept-possession.[99]

Bengson and Moffet consider their account to be a form of intellectualism, since concepts are clearly intellectual and ability simply stems from them. However, they give us no reason to accept that such concepts are simply constituted by the propositional rules that govern their

[99] John Bengson and Mark Moffet, " Know-How and Concept-Possession," *Philosophical Studies* 136 (2007): 390.

application. And this is what is necessary for their account to be a legitimate form of intellectualism. In fact, it seems that if conceptual mastery is to account for abilities, then concept possession requires more than knowledge of the necessary and sufficient conditions delineating a concept. Think of Mary, who doesn't have the ability to discriminate red, having never seen it, but certainly possesses reasonable conceptual mastery over the concept "red."[100]

The fact remains that an account of concepts is not an account of abilities. Reasonable conceptual mastery simply does not guarantee an explanation of how it is that knowledge is applied in action. After all, I may have the concept of a balance beam and the concept of a standing layout, and the concept of my body, but this in no way entails my body's ability to perform a standing layout on beam. And no amount of conceptual mastery, I don't think, could guarantee this ability.

At base, it seems that Bengson and Moffet have a nuanced definition of "concept." Their nuanced definition does not reduce knowing-how to knowing-that, but rather, it incorporates knowing-how into their requirements for concept-possession. This move is similar to Stanley and Williamson's "practical mode of presentation." For Bengson and Moffet, applying a concept appropriately is simply built into possessing a concept, just like for Stanley and Williamson, knowing that *w* is a way to do δ under a practical mode of presentation just is what accounts for being able to ϕ. Neither account explains how possessing propositional knowledge suffices for its successful application. Instead, both accounts conceal the fact that such an account is left wanting.

Further, for Bengson and Moffet, concept-possession must often be inferred from ability rather than the other way around. It is not the case that we will always find out that *S* possesses

[100] See Frank Jackson, "What Mary Didn't Know," *Journal of Philosophy* 83 (1986): 291-5.

concept *c*, and then infer that *S* is able to do *a;* rather, we say that because *S* is able to *a*, we know that *S* possesses concept *c.* Bengson and Moffet's ad hoc posit of a concept wherever ability is present no more entails the intellectualist legend, than observing order in the universe entails an intelligent designer.

Still, I do not claim that it is impossible that for some small group of purely cognitive tasks, a basic understanding of the rules that govern their application entails the ability to apply them intelligently.[101] But from the possibility that the concept of multiplication may entail the ability to multiply, it hardly follows that abilities are concepts. The take away point is this: almost always, propositional knowledge is not sufficient for ability.

4. 2 Differences in expertise are not proportional to differences in propositional knowledge

In the above gymnastics example, I posit that Mary Lou has less propositional knowledge about how to do a standing layout on beam than Bela, but she still has the ability to perform the skill while Bela does not. This suggests that expertise in performance is not proportional to differences in propositional knowledge. It is not always the case that the more one knows about an ability, the better one will be at performing it.[102]

In fact, quite frequently the opposite appears to be true. It is often noted that skilled performers regularly fail to describe the rules that govern their performance accurately. Consistently, experts either fail to state or state falsely, the rules that govern their behavior and the cues that they rely on in order to apply these rules.[103] That is, many skilled performers lack

[101] I think the most likely type of counterexample will follow from basic instances of math and logic (notice, how these differ from other abilities in being formal). But the fact that the way one learns these is not by memorizing rules, but rather by practicing their application, makes me think that even these will not get us very far.

[102] Being a linguist, after all, does not entail being a great speaker. And many people speak grammatically without being able to state the rules of grammar (anyone who has tried to teach ESL can attest to this).

[103] Various examples of this abound in the literature. See Charles Wallis, "Consciousness, Content and Know-how," *Synthese* 160 (2008): 130; Diane Berry and Donald Broadbent, "On the Relationship Between Task Performance and Associated Verbalizable Knowledge," *Quarterly Journal of Experimental Psychology: A Human Experimental Psychology* 36A (1984): 209-31; B. Brehmer, R. Hagafors and R. Johansson, "Cognitive Skills in Judgment:

the ability to report on the rules and facts that they actually use to successfully instantiate their abilities. This makes it unlikely that they instantiate their expertise following anything like the directives of a practical syllogism. After all, even if at the time of performance we do not require the conscious contemplation of a proposition in order to attribute propositional knowledge to an agent, it seems odd that someone in possession of the requisite conceptual repertoire, when prompted, could not reflect on and report her knowledge. If propositional knowledge was governing behavior, what kind of block could these experts have in accessing it? And why doesn't this happen with all sorts of other propositional beliefs?

Think for a moment about your own reasoning skills. Is it possible that even with knowledge of different logics, you still cannot cite the rules governing your comprehension of these sentences and the way in which you are coming up with objections to them? Try it. Compare this to reasoning through a problem using formal logic. Wouldn't it be easy for you to state the rules that you are following in that case? It seems that the difference is that in the second instance there are rules that you know and are actually following, while in the first, even if some rule exists that can be described as the rule that you are following, it is not the case that you know that rule and apply it appropriately in order to reason.[104] After all, the fact that behavior can be described as instantiating a rule does not entail that the agent knows the rule and performs the skill by applying it. I can bounce my leg up and down while listening to music and I'm sure that there's a description of my leg movements that can be given in rule form. Still, it is

---

Subjects' Ability to Use Information About Weights, Function Forms, and Organizing Principles," *Organizational Behavior and Human Performance* 26 (1980): 373-85; Arthur Reber & Selma Lewis, "Implicit Learning: An analysis of the Form and Structure of a Body of Tacit Knowledge," *Cognition* 5 (1977): 333-61.

[104] One may claim that the difference consists in the fact that in the first instance the rule is consciously contemplated while in the second it is not. This is hardly necessary, though. We can easily imagine a practiced logician not doing anything like consciously contemplating a rule that she is applying, but still being able to cite it as the rule that she is applying, when prompted.

unlikely that in order for me to keep rhythm with my leg, I need to know and apply anything even remotely resembling that rule.

This is especially obvious in cases of small children and animals. There are all sorts of creatures to whom we cannot reasonably attribute the type of sophisticated propositional states that would be necessary for them to possess, if abilities were really just discursive thoughts. Dogs explore their environments and play games and respond to familiar faces without having many of the concepts that are required for those abilities to follow from the application of propositional knowledge. A child may successfully reach for her bottle and smile at her mother, all the while lacking the concepts of "bottle," "milk," "smile," and "mother."

As Williams states,

> Shep the sheepdog has been successfully trained to round up sheep, many of us would feel comfortable in saying not only that it has acquired the ability to roundup sheep but also that it now know how to round up sheep…But no subject can hold beliefs that embody concepts which that subject fails to have, and Shep certainly does not have the concept of entropy. However it might have the concept of sheep.[105]

These examples present instances where ability exists in the absence of the corresponding propositional knowledge. These are instances where propositional knowledge is not necessary for knowing-how. This is important since if knowing-how were simply an instance of knowing-that, then we should not be able to have one without the other.

This is not to say that propositional knowledge is not necessary for any ability. Some abilities will most certainly require propositional knowledge for their successful instantiation. And it seems plausible that the more complex the ability, the more propositional knowledge will be required in order to perform it. But, for example, even if I have to know some things about

[105] Williams,108.

pianos and keys and timing and notes in order to play the piano, it is not the case that my ability simply follows from those facts. Remember Mary Lou and Bela: the differences in their abilities were inversely proportional to their propositional knowledge. This should never be the case if ability just is the possession of propositional knowledge.

Moreover, if knowledge-how were simply an instance of knowledge-that, then we should not expect contemplating propositional knowledge to interfere with the successful application of that knowledge in practice. But experts regularly claim that being "in the zone" requires *not* consciously reflecting on the rules governing their actions. [106] However, if ability just is thinking propositional thoughts, then why should doing it consciously interfere with success? After all, consciously contemplating propositions does not interfere with discursive thought.

The most likely explanation of why consciously contemplating a proposition interferes with skillful activity is because performing a skill requires a subject to focus on things other than propositional rules and facts. However, this simply reinforces the distinction between knowing-how and knowing-that—it reinforces the fact that whatever a subject must focus on is not propositional. But when thinking about propositions, we focus on propositions. As such, experts "in the zone" uphold the disparity between purely intellectual thought and its practical instantiation.

Crucially, what we've seen in the previous sections is that (1) Ryle's regress stands, (2) it is possible to have knowing-that without knowing-how, and (3) it is possible to have knowing how without knowing-that. *Ryle's regress plus the double dissociation of knowing-how and knowing-that should make us certain that these two categories are not identical.*

5. What knowing-how is not

[106] This point is made by Hubert L. Dreyfus, "The Return of the Myth of the Mental," *Inquiry* 50 (2007): 352-65.

Though, I have argued that knowing-how cannot survive a straightforward reduction to knowing-that, I have not said what it is that makes knowing-how different from knowing-that. In this section, I will briefly comment on what knowing-how is *not*.

5. 1 Know-how is not theoretical or practical reasoning

It is almost universally accepted that know-how is not theoretical knowledge or reasoning. This is because knowledge-how applies to the realm of action and practice, and theoretical reasoning does not.[107] Though theoretical reasoning can produce deductively valid conclusions, these conclusions are not instructions for action. And while theoretical knowledge can tell us what is true, it cannot tell us how to translate this truth into successful practice. For example, theoretical reasoning can tell us that (1) All horses are animals, that (2) all animals need food to survive, and so, (3) all horses need food to survive, but it cannot tell us that we ought to feed the horses. According to Aristotle, these are the responsibilities of practical reason.[108]

As opposed to theoretical reason, practical reason is reasoning about action. It is instrumental, means-ends reasoning about how to reach one's goals.[109] As such, it seems much more likely that practical reasoning will be able to account for know-how. After all, if practical reason can tell us what to do given particular desires, and know-how is the ability to do things, then at least we are playing the same game here. Practical reason, for instance, can tell us that: (1) When one wants to run a marathon, one must train, and, since (2) I want to run a marathon it follows that, (3) I must train. Without some obstacle standing in my way, practical reason should spur my training.

---

[107] See Aristotle, *Nicomachean Ethics*, Book VI.
[108] Also, recall the is/ought distinction in David Hume, *Treatise of Human Nature* (USA: Prometheus Books, 1992).
[109] See Aristotle, *Nicomachean Ethics*, Book VI.

Importantly, the conclusion of this practical syllogism is not an instruction on how to train, but simply the imperative to do so. But we need more than this to account for know-how; we want to know both *that* I ought to train and *how* I ought to train myself, in the middle of winter, with only three months left before the marathon and with this stiff knee, to boot. One may think that if we just get more specific then we can get the practical syllogism to account for the nuances of successful action in particular circumstances. But the *how* of the practical syllogism is too weak to act as a guide for implementing behavior. After all, successfully performed ability must be responsive to the subtleties and nuances of the exact situation in which one is performing. But, the universal or major premise of the practical syllogism cannot apply to particular situations. If it did, it would cease to be a universal premise.

The more specific one is with the details of the major premise of a practical syllogism, the less it seems one's reasoning is going from the general to the particular. But the less particular one is with one's reasoning, the less likely it is that that reasoning can be responsible for one's ability to successfully perform a skill. Because actions must be executed in particular settings with an acute sensitivity and responsiveness to the actual nuances of this very situation, here and now, the major premise of a practical syllogism would have to incorporate all of the very particular details of this very situation, here and now, in order to be responsible for guiding action. Yet, this sort of detail undermines the purpose of having general rules from which to derive practical directives.

For example, let us say that I would like practical reasoning to direct me in playing a note on the piano. We'd have to start with something like, "When one wants to play a note on the piano, then one ought to depress a piano key." But this is not enough to guide behavior, for we need to know what to depress the piano key with, and so, we need to also say that "When one

want to depress a piano key, one ought to use one's finger." But then we have to say what part of one's finger one should use to depress the key, and so we have to say that "when one wants to use one's finger to depress a piano key, one ought to use one's finger tip, at such and such an angle." But then we should want to know how hard to press, and so, we'd have to say that "when one wants to depress a piano key with one's finger tip at such and such an angle, then one must press down with the appropriate amount of pressure." But surely we need to know what the appropriate amount of pressure is, and in order to find this out we must ask what kind of piano I am playing, and also we should find out when it was tuned last. And from this, we'd have to construct another more detailed premise. But this still isn't enough, for what if I want to play the note loudly? Then we must ask how big the room is, and its shape, and we must find out what is covering the walls. But surely all of this cannot be incorporated into the practical syllogism. Or rather, even if it could be, what ends up happening is that the practical syllogism doesn't seem to do the work that it is meant to do. This is because to reason by way of the practical syllogism, one must have general rules for how to behave in particular situations. But as it turns out, if one is to use practical reason in order to account for skill, then one does not have general rules, but only particular ones. And problematically, particular rules are not rules at all.

In order to respond appropriately to any particular situation one must take into consideration much more than general rules. The specificity of action makes it impossible that such rules are sufficient for directing and guiding behavior.[110]

[110] When thinking about this problem, I cannot help but be reminded of the movie *Borat*, Directed by Larry Charles (New York: 20th Century Fox, 2006) where Sasha-Baron Cohen is being coached on how to make a not-joke. The multiplicity of ways that Cohen misinterprets the instructions just goes to show how nuanced our understanding and sensitivity must be in order to perform even a seemingly simple action successfully. Mr. Bean is another good example of how much can go wrong, and so, how much we take for-granted what is necessary for something to go right.

Aristotle admits this much in Book II of the *Ethics* where he—in no uncertain terms—claims that determining the mean of ethical action is a matter of perception and not reasoning. He states:

> But up to what point and to what extent a man must deviate before he becomes blameworthy is not easy to determine by reasoning, any more than anything else that is perceived by the senses; such things depend on particular facts, and the decision rests with perception.[111]

Aristotle is clear that in particular disciplines like ethics, medicine, and navigation, the appropriate action will not be the result of applying a rule. He argues that in circumstances that require particular consideration of situational facts, general reasoning lacks the power to provide a course of action. This is because such reasoning is incapable of taking into account the innumerable variables that require consideration in order to generate an appropriate response. As such, know-how cannot be the result of practical reasoning.

5.2 Know-how is not conceptual

Further, since skills are concerned with the realm of action and action is always in a particular situation, on a particular material, in a particular environment, we should not expect skills to be constituted by concepts. After all, at the very heart of conceptuality is that idea that concepts can be abstracted away from particular circumstances in order to enter into various intentional relationships while retaining their identity. Concepts are, by definition, general and recombinatorial.[112] But abstracting away from a situation is devastating to the instantiation of a skill, since it undermines the successful performance of it: in order for a skill to be successful it must adjust, shift and respond to the very particular circumstances of the environment in which it is being instantiated.

---

[111] Aristotle, *Nicomachean Ethics*, Book II, 1109a18.
[112] See Gareth Evans, *The Varieties of Reference* (Oxford: Oxford University Press, 1982).

So, while concepts can be abstracted away from their multifarious environments, skills develop by becoming more and more attuned to their particular circumstances. Though skills can be instantiated at different times and places and, in this sense, may be recombinatorial, it does not make sense to think of a skill existing "as such." There is no skill outside of the particular instantiation and environment in which it becomes manifest. As such, skills are uniquely embedded in their particularity. While concepts can be considered generally, the procedures of a skill lack this non-worldly quality. They are adaptable precisely because they are not abstract, but rather, sensitive and responsive to particular situations.

Of course, there must be a way for us to say that the same skill is instantiated in various circumstances and this may make it seem that a skill, like a concept, can retain its identity despite a multiplicity of conditions. But the fact remains that being applied in different situations does not entail being a concept. To be a concept requires being both recombinatorial and satisfying the generality constraint. Quite clearly, skills cannot do the latter.

5.3 Know-how is not propositional:

From the fact that skills are not conceptual, it follows that skills cannot be propositional. This is because propositions are necessarily constituted by concepts. So, though skills may be related to facts, rules, theories, and all sorts of propositions, they are not identical to any of them. For all such states must be composed of concepts and concepts are incapable of being particular in the way skills require. Since skill is utterly embedded in an environment, its success is dependent on the capacity to adjust and respond to the very circumstances in which it is located. That is, even if there are general rules governing the way to successfully perform an action (and of course there are), such general rules can never provide the necessary directives for acting *en situ.* General rules or facts are often quite beside the point, since acting *en situ* requires not the

application of universal truths, but the discernment of nuances specific to the context in which one is acting. This is what skill must be and this is what concepts and propositions are not.

5.4 Know-how is not aimed at truth

Because skill deals with the realm of action, rather than with the realm of fact, it has been largely accepted that knowing-how is not aimed at truth, but rather at success.[113] So, while getting propositional states right means getting them to hit the mark of truth, getting knowing-how right means successfully achieving one's end. This is an important aspect of skill, since it not only provides us with a direction or a goal, but also gains us entrance into the realm of normativity—of right and wrong, good and bad, success and failure. Though skill is not measured in terms of truth, this does not entail that it must be reduced to a non-normative phenomenon.

The fact that skills are not true or false is quite significant since it implies that skill cannot enter into logical argument and reasoning. Since the goal of formal reasoning is to preserve truth under logical transformation, we should not expect that skills can enter into the inferential relations necessary for such reasoning. After all, if they are not true in the first place, then there is no sense in trying to preserve their truth under logical transformation. And if they cannot enter into logical deduction, then it seems that at the very least we have a reason for questioning whether they are identical to any conclusion that follows from it.

To summarize, whatever know-how turns out to be, it is not theoretical reasoning, it is not practical reasoning, it is not conceptual, it is not propositional and it does not aim at truth. The next challenge, of course, is to provide a positive account of know-how, rather than just saying what it is not.

---

[113] See Carr, "Knowledge in Practice;" Michael Dummett, *The Origins of Analytic Philosophy* (Cambridge, MA: Harvard University Press, 1993) and Susan Hurley, "Making Sense of Animals," in *Rational Animals?* ed. Susan Hurley and Matthew Nudds, 139-171 (Oxford: Oxford University Press, 2006).

I take it that this chapter, by presenting a comprehensive defense of the distinction between knowing-how and knowing-that, constitutes a crucial step in laying the groundwork to make such an analysis possible. It is crucial to have this foundation from whence to proceed with an investigation into the type of unique knowledge that is operative in embodied skill. The defense of the distinction between knowing-how and knowing-that secures the legitimacy and mandate necessary to consider alternate types of knowledge, not simply as derivative of propositional thought, but as requiring unique and substantive explanations of their own.

CHAPTER 4

1. An introduction to a positive account of skill

In this chapter, I attempt to define the category of skill. My overall goal is to argue that if skills hold logical or semantic relationships to the processing of early perceptual systems, then those systems are cognitively penetrable. This is because skills, as I will argue, are characterized by practical intelligence. From hereon, I will use the terms "skill," "intelligent ability," and "know-how" interchangeably.

A satisfying definition of skill should offer a non-arbitrary distinction between playing the violin and salivating, and also between the ability to read and the knowledge that letters represent sounds. As such, to begin homing the category of skills, I will seek to avoid two extremes. The first is to differentiate skills from tropistic, instinctive, reflexive and basically unintelligent behaviors, and the second is to avoid identifying skill with propositional thought. Skill should turn out to be intelligent, but not intelligent in the same way that truth-functional, discursive states are intelligent. Skills should exhibit paradigmatically cognitive characteristics that are missing from unintelligent behaviors, but that are not reducible to the knowledge of facts, either.

It should be quite obvious that skills are closely allied with abilities. However, they are not, pace Gilbert Ryle, identical to them.[114] The category of skill is smaller than the category of ability, though the two are intimately related.[115] I claim that, though some abilities are not skills, all skills are abilities. We can think of the relationship between the two as one of genus and

[114] David Carr, "The Logic of Know How and Ability," *Mind* 88 (1979): 409 has argued this point: namely, that know-how is intentional ability.

[115] This is because the second criterion of skill is the criterion for abilities. This shared criterion should be sufficient to elucidate the relationship between skill and ability.

species. This should be quite clear since, for example, my lawn mower has the ability to cut grass (it would be pretty useless if it didn't), but my lawn mower does not know how to cut grass; it does not possess the skill of grass cutting. On the other hand, every time I know how to do something, I have the ability to do it, though of course, I may lack the opportunity.[116] If I didn't have the ability to *a*, it would be quite mysterious as to what it might mean to say that I had the skill of *a*-ing. So, everything I know how to do, I am able to do; but not everything I am able to do is something that I know how to do.

Another consideration for defining the category of skill is to both connect and differentiate skills from intentional actions. As is the case with abilities, intentional actions and skill instantiations hold a relationship of genus to species. That is, though all skills, when they are instantiated, are intentional actions, not all intentional actions are instantiations of skills.

After all, it would seem patently odd if any of the following examples of intentional action qualified as instantiations of a skill:

(i) The last time I played basketball was in my high school P.E. class. Today, I am asked to try to make a basket with my back towards the hoop and my eyes closed. Incredibly, I make it.

(ii) I have an itch, so I scratch it.

(iii) For months, I have tried to see those 3-D pictures in 3-D. I try, but fail.

From these examples it ought to be clear that skill instantiations are different from these types of intentional actions: (1) those that are executed intentionally, but are only successful as a result of dumb luck or accident, (2) successful actions that are intentional, non-accidental, but do not exhibit the control or sophistication characteristic of skill, and (3) intentional actions that end in failure.

---

[116] That is, having a skill does not mean being able to successfully perform it under *any* conditions. This point will become clearer when I discuss criterion 2 and why we qualify it with "in conditions *c*."

As such, the instantiation of a skill will reside in the overlap between intentional action and ability. But skill requires something more than the satisfaction of these two criteria. Being done intentionally and successfully are necessary but not sufficient conditions for some event *e* to count as an instantiation of a skill. This must be so, since example (ii) is both an intentional action and an ability, but it is not in any obvious way a knowing-how or a skill. I will posit that for some event *e* to legitimately count as a skill, such an event will (1) have to be instantiated intentionally, (2) one will have the ability to instantiate it successfully, but also, (3) it must be characterized by practical intelligence—or as I will argue, it must be developed as a result of practice; that very special type of learning.

Using the above considerations as bumpers, in the next section, I will present three definitive criteria of skill. Before so doing, however, I'd like to defend my choice to focus on specifically embodied skills in the remaining chapters. I will focus on skills that are expressed in physical activities, and not merely in mental operations, because it seems plausible that perceptual processing will most likely be impacted by events that change the relationship that a subject holds to her environment. Since perception is always of the world from a particular perspective, then, because embodied actions determine a subject's location in and interaction with the world, they are the best candidates for the kinds of actions that could impact perceptual processing. This does not imply that if purely cognitive processes like doing math or composing music in one's head were to affect perceptual processing, then such interactions would not count as instances of cognitive penetration. They certainly would. But, for the reasons stated above, I will be careful to focus on embodied skills, since they mediate interaction between subject and environment.

2. Skill: a definition

I define skill according to the following criteria:

(1) If *a* is a skill of *S*'s, then when *S* instantiates *a*, *S* must have some reason *r* for *a*-ing, and *r* must be responsible for *S*'s instantiation of *a*.

(2) If *a* is a skill of *S*'s, then under conditions *c*, *S* is usually able to instantiate *a* successfully.

(3) If *a* is a skill of *S*'s at *t*, then there was some time *t*-1 that *S* could not instantiate *a* successfully, and *S* has learned to instantiate *a* successfully as a result of practice.

3. Criterion 1: intentional action

Criterion (1) requires that skills are performed for reasons. This means that for some event *e* to qualify as an instantiation of skill *a*, that event must be an intentional action. An action is intentional when a subject performs it with some end in mind and because the action is believed to bring one closer to achieving that end. Importantly, an intentional action is not something that just happens to one.[117] The subject of an intentional action is more than a subject—she is an agent.[118]

This requirement for skill is quite uncontroversial because it makes a fairly straightforward demand: when one knows how to do something, one actually has to *do* (to be the agent of) the thing that she knows how to do. Or: if *a* is a skill of *S*'s then *S* must be the agent of *a*'s instantiation and not just the recipient of *a*'s happening. As the Jason Stanley and Timothy Williamson discussion made clear,[119] not all things that people can be said to do are things that they know how to do. Hannah digests her food, but she does not know how to digest her food. She is not the agent of the action, though she is its subject: digestion happens inside of

[117] We must differentiate, after all, between being the subject of an intentional action, such as being hit in the face with a good right hook, and being its agent.
[118] This issue is widely discussed in debates about moral responsibility and free will. See especially, Harry Frankfurt, "The Problem of Action," *American Philosophical Quarterly* 15 (1978): 157-62 and Donald Davidson, *Essays on Actions and Events* (Oxford: Oxford University Press, 2001), 43-62.
[119] See chapter three, section 3.

Hannah.[120] For this reason, digesting food cannot be something that Hannah has the skill of doing.

The main function of this first criterion is to differentiate skills from unintentional happenings, random events and also from systematic, teleological but nonintentional processes. Acting for reasons requires that we make these distinctions because doing something for a reason is not only different from something happening for no reason,[121] but it is also different from goal-oriented processes that possess direction, but lack agency. [122] For instance, phototropism is a teleological process: plants move towards the sun because light is required for photosynthesis and plants need to undergo photosynthesis in order to survive. The plant behaves in a way that may be described as reasonable, but it does not behave in this way *because* of reasons; they are not reasons for it.[123] As such, the plant does not act intentionally.

3.1 Desire and belief

Formally, an action is done for a reason when a pro attitude and an instrumental belief combine to cause the action.[124] As Donald Davidson claims,

> Whenever someone does something for a reason…he can be characterized as (a) having some sort of pro attitude toward actions of a certain kind, and (b) believing (or knowing, perceiving, noticing, remembering) this action is of that kind.[125]

[120] See Alva Noë, "Against Intellectualism," *Analysis* 65 (2005): 279.
[121] This does not mean that such an event is uncaused, but just that there is no obviously coherent reason that can account for the happening. Hence, "random." E.g., the car randomly broke down. There is clearly a cause for the car breaking down, but it is not related to the behavior of the car in such a way that we can give a coherent account of the car's behavior and the breaking down of the car together.
[122] See Fred Dretske, *Explaining Behavior* (Cambridge, MA: MIT Press, 1988), 109-122 for the distinction between goal-directed and goal-intended behavior.
[123] This is the way Fred Dretske, "Minimal Rationality," in *Rational Animals?* ed. Susan Hurley and Mathew Nudds, 107-115 (Oxford: Oxford University Press, 2006). differentiates minimally rational behavior from others. See also chapter two, section 3.2 and this chapter, section 3.2 for a discussion of minimal rationality.
[124] I will follow Davidson and Dretske, *Explaining Behavior,* fairly closely here.
[125] Davidson, 3-4.

A pro attitude is to be understood broadly as a desire. As Davidson explains, the category of pro attitudes includes,

> [D]esires, wantings, urges, promptings, and a great variety of moral views, aesthetic principles, economic prejudices, social conventions, and public and private goals and values in so far as these can be interpreted as attitudes of an agent directed toward actions of a certain kind.[126]

Very generally, in order for an action to be done for a reason, the acting agent must desire something toward which the action is directed. The action must be goal-directed. So, if an agent is hungry and orders food at a restaurant, the wanting of the food is that part of the agent's reason, which comprises the pro attitude.

The second element, which is necessary for an action to be done for reasons, is belief. As Fred Dretske has argued, in order to connect some movement to a desire, it is not enough that an agent has a desire, but it is also required that the agent performs the movement in order to satisfy the desire.[127] So, specifically, an agent must have an instrumental belief holding that the movement is a means for achieving the desired end. In the above example, if an agent acts for reasons when he orders food in a restaurant, then he must not only be hungry, i.e., have a pro attitude towards food, but he must also believe that ordering food, amongst other things, will lead to the satiation of his hunger.

3.2 The limits of intentional action

On Dretske's account, in order for behavior to be minimally rational, or for it to be done for reasons, it must be that the behavior is connected, in the right way, to intentional states.[128] Importantly, in order for a behavior to be done for reasons, the content of an intentional state

---

[126] ibid, 4.

[127] "[T]o explain someone's performance it isn't enough to point out that, in the conditions he believes to obtain, his behavior leads to results he desires. One also wants to know whether the agent *knows* (or at least believes) that his behavior *will* lead to results he desires Dretske, *Explaining Behavior*, 116.

[128] It seems reasonable to expect that intentional actions are minimally rational.

must be explanatorily significant and not simply causally efficacious in causing the behavior.[129] So, when we look at acting for reasons, it is not enough that one can attribute representational states to an agent. Rather, it must be the semantic content of those states that explains the behavior. This is important because this requirement will limit the number of actions that are done for reasons and, thus, the number of actions that will end up satisfying the first criterion of skill.

So, while we can say that a plant has representational states indicating the location of the sun, it is not the case that the semantic content of those states explains the plant's movements. Its behavior is reasonable, but not done for reasons. As Dretske states, "something possessing content or having meaning, can be a cause without its possessing that content or having that meaning being at all relevant to its causal powers".[130] In order for it to be a reason and not just a cause, it is required that the meaningful content of the intentional state explain the behavior.

Significantly, on Dretske's account, learning and the ability to respond to the content of an intentional state is integrally linked to acting for reasons. He states,

> the reason learning is so central to *intelligent* behavior, to the behavior of *people,* is that learning is the process in which internal indicators… are harnessed to output and thus become relevant—as representations, as *reasons*—to the explanation of behavior of which they are a part. It is in the learning process that information-carrying elements get a job to do *because* of the information they carry and hence acquire, by means of their *content,* a role in the explanation of behavior.[131]

In order to differentiate intentional behavior from behavior of other kinds, we can ask if a change in the meaning of an intentional state could lead to a change in behavior. Is it possible for the intentional state to mean anything other than what it does? After all, if a subject is acting for reasons, then a change in the reason should be reflected in a change in behavior.

---

[129] Recall the opera singer whose singing shatters a glass. See chapter two, section 3.2.
[130] Dretske, *Explaining Behavior,* 80.
[131] Dretske, *Explaining Behavior,* 104.

The reason that we cannot claim that the semantic content of the plant's informational states explains its behavior lies partly in the fact that the plant's informational states cannot represent anything other than what they do. That is, even if a plant were no longer to derive nourishment from the sun, the plant could not change the fact that it represents the sun as a source of nourishment. The plant cannot act for different reasons and so it seems unlikely that it is acting for any reasons at all.[132] In order for there to be reasons for behavior, those reasons must possess a degree of flexibility in representing the significance of environmental conditions. While indicator functions may be fixed, reasons are not.[133]

Daniel Dennett, pace Dretske, has argued that the behaviors of plants (and other organisms of this sort) do in fact meet the criteria for minimal rationality. He argues that learning in a lifetime versus learning over generations is an arbitrary way of marking out the relevant distinction concerning the rationality of behavior. Dennett claims that plants act for reasons, and though their reasons change by way of evolution rather than personal development, this does not undermine their being reasons for action.[134] He argues that because indicator functions can change through evolution and influence behavior, such behaviors, since they hold logical relations to shifting environmental conditions, should be considered intentional. Notice that Dennett agrees with Dretske's claim that representational states must be able to change in meaning in order for them to function as reasons. Dennet and Dretske disagree, however, on the

---

[132] I think of this point as in the spirit of Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell Publishing, 1953), 84 and Crispin Wright, "Wittgenstein's Later Philosophy of Mind: Sensation, Privacy, and Intention," *Journal of Philosophy* 86 (1989): 634 on introspection's infallibility. Wittgenstein argues that it makes no sense to say of instrospective states that they are always true, because for "true" to mean anything, the possibility of error is required.

[133] See Fred Dretske, *Naturalizing the Mind* (Cambridge, MA: MIT Press, 1997), 8-22 for a discussion of the relationship between having the function of indicating and representational content

[134] Daniel Dennett, "Ways of Establishing Harmony," in *Dretske and His Critics*, ed. Brian P. McLaughlin, 118-130 (Cambridge, MA: Blackwell, 1991).

amount of flexibility that a representation must possess in order for it to count as a reason. They disagree on the thresholds of intentional behavior, but not on the requirements.

For my purposes, I will side with Dretske on this issue. That is, I will take a conservative stance as to which actions are performed for reasons. Regardless of my commitment, however, we ought to notice two things: (1) if Dennett is right then we must expand the category of intentional actions, and thus, potentially, but not necessarily, expand the category of skills, too. That is, siding with Dennett does not undermine my claim that skills are a species of intentional action, rather, it strengthens it. Siding with Dennett simply broadens the category of actions that may turn out to be instantiations of skills. Further, (2) even if Dennett is right and the behaviors that result from the accumulation of learning over generations count as intentional and rational, we can still make a distinction between intentional actions that are flexible as a result of learning in a lifetime and those that require learning over many generations. Given that we can make this distinction, an analysis of skill can effectively deal exclusively with those kinds of intentional actions that are susceptible to intra-lifetime changes. Since these are more intuitive instances of acting for reasons, such a move should hardly require an extensive defense. In fact, it is not necessary for me to settle the dispute about whether plant behavior is intentional or not. I can simply say the following: there are clear cases of intentional action, and then there are less clear cases. When thinking about skill, I will focus on the cases that are clearest.[135]

3.3 Reasons and causes

[135] A similar kind of scope issue will come up when I discuss the practice in chapter five. There too, for the same reasons, I will circumscribe the actions that count as practice conservatively.

A problem that has traditionally been considered central to the issue of intentional action is whether it is possible for reasons to be causes.[136] Rather than engage in a lengthy debate about this problem, I will take what I find to be the most minimal and reasonable solution. Following Michael Bratman, I will consider reasons to be the causes of actions in a functionalist sense.[137] That is, reasons are states that tend to promote particular predictable responses in a given context. So, though I do not claim that reasons are causes in any necessary, law-like sense, [138] a functionalist approach elucidates how reasons are related to actions in ways that are explanatorily significant and can serve as the basis for generally reliable predictions.

It is vital, however, that I emphasize that the relationship between reasons and actions must be a causal one, because one may have a reason for doing something and do it, but not do it as the result of her reasons. For example, Andrew may have a pro attitude towards food, and believe that ordering food at a restaurant is a way for him to satisfy his hunger, and further, he may order food, and yet his ordering food may still not be done for reasons. Imagine that Andrew is hungry and wants to order delivery and plans to order delivery after checking his email. But after turning on his computer, he accidentally presses the wrong key and completes an online delivery request. In this situation, Andrew has reasons for ordering food, and he orders it, but he does not order it for reasons. This is because the reasons that he has for ordering food were not responsible for his actually ordering the food. What we need, then, in order for an action to be performed for reasons, is for a reason to act as the functional cause of the action.

3.4 Individuating skill instantiations and individuating skills

---

[136] See Davidson, *Essays on Actions and Events*, 3-20; Elizabeth Anscombe, *Intention* (Cambridge, MA: Harvard University Press, 1957), 15-18; Gilbert Ryle, *The Concept of Mind* (Chicago: The University of Chicago Press, 1949), 62-82.
[137] Michael Bratman, *Intention, Plans, and Practical Reason* (Cambridge, MA: Harvard University Press, 1987).
[138] See David Hume, *Treatise of Human Nature* (New York: Prometheus Books, 1992), 78-82 for a discussion of causation and necessity.

A natural response to criterion (1) may be to assume that if actions are performed for reasons, then the types of actions that they are should be individuated according to the content of the reasons one has for performing them. So, if I order food at a restaurant because I am hungry and believe that ordering food will result in my being able to eat it and thus satiate my hunger, then the act of ordering food should be individuated according to what I believe it to be: ordering food because I'm hungry and believe that ordering food will result in the satiation of my hunger. Individuating actions in this way, however, leads to some thoroughly unpalatable consequences.

For instance, mistakes become impossible, if we individuate intentional actions according to the reasons one has for performing them. This is because by failing to correspond to a reason, an action will either become unintentional (i.e., not something that an agent does for reasons) or even more bizarrely, a successful instance of the intended action.

Consider this example: if I have a reason for soft-boiling an egg, but leave the egg in boiling water too long and it ends up hard-boiled, then it seems that we'd have to claim that the action of hard-boiling an egg was not an action that was performed with intentions. This follows since I never had a reason for hard-boiling an egg and if intentional actions are individuated according to the reasons one has for performing them, then I never had an intention for this action. As such, it becomes the case that hard-boiling an egg was not something that I did, but rather, something that happened to me, or something I did voluntarily but for no reason. But this is an odd consequence, given that I acted for reasons in preparing the conditions for hard-boiling the egg.[139]

[139] Questions of intentional action are often immediately associated with issues of responsibility. To make a stronger case, we can think of a hunter hunting in the woods, firing a gun at a buck, but shooting a hiker instead. It would be strange if we claimed that the hunter's shooting the hiker wasn't something that he did, but rather, something that just happened. It seems much more likely to say that the hunter shot the hiker (performed the intentional action of shooting the gun), but did not intend to shoot the hiker (the consequences were unintended). See also Davidson, *Essays on Actions and Events*, 45.

The problem with individuating intentional actions by way of the reasons one has for performing them is that we commit ourselves to a view where either all unintended consequences result from actions that just happen to one or are actions that are not caused by one's beliefs and desires. This sets up a theory where failed intentions undermine a robust sense of agency. This must be the case because there are no reasons that correspond to an action that fails to meet its intention, but a corresponding reason should be exactly what makes some event an intentional action of any particular kind.

Maybe even more strangely, it may turn out that my hard-boiling an egg is really an instance of my soft-boiling an egg. If actions are individuated according to the reasons one has for doing them, then this seems to follow. After all, my reason for performing the action of hard-boiling the egg, in fact, was to soft-boil the egg. But this is clearly unacceptable. The crucial point is this: any theory of individuating intentional actions must have the ability to do justice to both the reason for the action and the extensional description of the action. This is the only way that a theory can allow for the possibility of mistakes.[140]

As such, intentional actions cannot be individuated, in any straightforward way, according to the reasons one has for performing them. This does not, of course, mean that we should disregard the reasons one has for performing actions. It only means that those reasons cannot be the mark by which we identify the type of event of which the action is an instance.

Following Davidson, I suggest that rather than individuating actions according to the reasons one has for doing them, instead, we should individuate actions extensionally. A reason can be a criterion for intentional action without thereby becoming the mark by which that action or skill instantiation is identified. As Davidson writes, "If we can, as I am urging, that a person does, as

[140] See Alfred R. Mele and Paul K. Moser, "Intentional Action," *Noûs* 28 (1994): 39-68 for a similar point; Gilbert Harman, "Practical Reasoning," in The Philosophy of Action, ed. Alfred R. Mele (Oxford: Oxford University Press, 1997).

an agent whatever he does intentionally under some description, then, although the *criterion* of agency is, in the semantic sense, intentional, the *expression* of agency is itself purely extensional."[141] That is, one can require that intentional actions are done for reasons but not individuate those actions according to the content of the reasons for which they are done. And since acting for reasons is required for intentional action, we do not run into problems of behaviorism. After all, reasons for actions are taken into account on Davidson's theory, they are just not cited as the ways according to which actions are individuated.

Further, out of a desire for ontological conservatism, we can also side with Davidson in holding that descriptions of actions with identical extensional boundaries are descriptions of identical actions. Thus, different extensional descriptions of skill instantiations will not amount to instantiations of different skills, but merely different descriptions of the same skill. For example, if the Moonlight Sonata is identical to Beethoven's favorite composition, then if I have the skill of playing the Moonlight Sonata, I also have the skill of playing Beethoven's favorite composition. Using this approach to skill, we can individuate skills extensionally and avoid multiplying entities unnecessarily by abstracting away from disparate descriptions of identical skill instantiations.

When we return to the above example and individuate actions extensionally but require that there is some reason for them, we are in a position to say that I hard-boiled the egg intentionally, though I did not intend to hard-boil it. I can act for reasons, without guaranteeing that my action successfully reaches its goal. Further, if I describe the action first as hard-boiling an egg and a second time as making breakfast, it does not turn out that these are two descriptions are distinct actions. And if they meet the other two criteria of skillhood, they will not have to be identified as instantiations of two different skills.

---

[141] Davidson, *Essays on Actions and Events*, 47.

Though it is preferable not to individuate actions according to the reasons for which they are performed, we should still want successful intentional actions to correspond to the reasons for which they are instantiated. After all, we want to avoid a theory of skill where one's reasons for acting are completely irrelevant to the skill performed. So, even if we do not individuate intentional actions according to reasons, we should still want some sort of correspondence or matching between reasons and the resultant instantiations of skills, when those instantiations are successful.

I propose that in order for an action to be an instantiation of a skill and to meet criterion (1), then the reason one has for performing *a* must match at least one extensional description of *a*. This is because, though skills are not individuated according to the reasons one has for performing them, we must require that the reason that causes *a* is really is a reason for *a*-ing under some description, rather than a reason for doing anything whatsoever. This entails that there must be some correspondence between one's reason for *a*-ing and the actual activity which is an instance of *a*-ing. Of course, this also implies that one may have a reason for playing the Moonlight Sonata, and since the Moonlight Sonata is, let's say, identical to Beethoven's favorite composition, one can intentionally play Beethoven's favorite composition without having any intentional content that mentions Beethoven's favorite composition at all.

This brings us to the issue of the opacity of reasons and whether or not they must translate into the opacity of skills. Everyone agrees that I can have a reason for playing the Moonlight Sonata but not have a reason for playing Beethoven's favorite composition, if I do not know that the Moonlight Sonata *is* Beethoven's favorite composition. However, if the two are identical, then I cannot play one without playing the other. On my proposed theory, I have the skill of playing both as long as I have a reason for playing either.

On my account, I may know how to play Beethoven's favorite composition without possessing a reason that explicitly mentions playing Beethoven's favorite composition, just as long as I have a reason for doing something extensionally identical to playing Beethoven's favorite composition. Here, it turns out that I may know how to play Beethoven's favorite composition without knowing that I know how to do this.

One may object, of course, that it's odd that we may have skills that we cannot perform when asked. After all, if you ask me to play Beethoven's favorite composition, I will be unable to do so if I do not know that Beethoven's favorite composition is identical to the Moonlight Sonata. If I do not know that I know how to play Beethoven's favorite composition, even if I am in fact able to play it, then I will not even be able to try to play it. It seems wrong to claim that I know how to do this thing that I cannot do, when asked.[142]

I admit that this is a seemingly disastrous consequence of individuating skills in the manner that I have proposed, but I insist that the disaster is only apparent. This becomes quite clear if only we ask what I need to do in order to be able to respond to the request to play Beethoven's favorite composition. What I need to learn is the fact that Beethoven's favorite composition is identical to the Moonlight Sonata. I do not need to learn to place my hands on the piano keys correctly; to strike the notes with the appropriate force; to learn to read notes; to practice this particular sonata, etc. But if I did not know how to play Beethoven's favorite composition/the Moonlight Sonata, this is quite clearly what I would have to learn. Further, it wouldn't be surprising if, when presented the sheet music or a sample recording, I exclaimed, "Oh! I know how to play *that*!" My point is this: I do learn something new when I learn that the

[142] Katherine Hawley, "Success and Knowledge How," *American Philosophical Quarterly* 40 (2003): 26-28 and John N.Williams, "Propositional Knowledge and Know-How," *Synthese* 165 (2008): section 7, among others, agree that the ability to try is key to having an intentional ability. I hope that I have provided good reasons to question whether being able to try is a prerequisite of knowing-how to *a*. See also David Carr, "The Logic of Knowing How and Ability," 408.

Moonlight Sonata is Beethoven's favorite composition. I learn the fact that they are identical. What I do not learn, however, is a new skill.

This means that, the opacity of reasons need not entail the opacity of skills.[143] It is perfectly natural to say that I am able to learn that I possess skills of which I am not aware. This position may seem a bit counterintuitive, but I insist that it should not make us uncomfortable. Why, after all, should we think that everyone that knows how to do *a*, must know that she knows how to do *a*? The KK thesis,[144] in all its dualist glory, has numerous unpalatable consequences: it rules out that animals and children know how to do things, it unduly emphasizes our introspective capacities, and it arguably begins a regress that makes any kind of knowledge impossible. For these reasons, it is best to avoid any sort of robust internalism, whether for propositional states or skills.[145]

The problem is that theorists have often taken the opacity of reasons to entail the opacity of skills. So, they argue that it is possible to know how to do *a*, but not *b*, even though *b* is identical to *a*. On my account, we can retain the opacity of reasons, and also require that skills are performed for some reason that matches at least one accurate description of *a*, but still say that if *S* knows how to do *a*, and *a* is identical to *b*, then *S* knows how to do *b*. This approach makes it easy to explain why it is that sometimes we can discover that we already have skills hidden within our repertoire.

---

[143] Pace, Carr: "so it appears that sentences about knowing how, unlike those about ability, are truly non-extensional," "The Logic of Knowing How and Ability," 408.

[144] The KK thesis (the knowledge of knowledge thesis) states that to know that *p*, one must know that (why) she knows that *p*. It requires that one can access the justification of what one truly believes. This internalist account of knowledge takes the ability to reflect on one's knowledge as a prerequisite for possessing it. See Stanford Encyclopedia of Philosophy, "Internalist vs. Externalist Conceptions of Epistemic Justification," http://plato.stanford.edu/entries/justep-intext/, section 1.

[145] The KK thesis causes a regress because to know how to do something, I must know why I know how to do it, and if I am to know that I know how to do it, then I must know why I know that I know how to do it, and so on, *ad infinitum.* Additionally, why should my beliefs about myself hold any special place in judging my abilities or my knowledge? After all, I don't need to know what you know how to do in order for you to know how to do it. And given that introspection, often goes wrong, it seems misguided to rely on such knowledge as a prerequisite for knowing-how.

Further, on this account we can explain why two people may possess the same skill, even though they describe it differently. Take for instance the following example: Mom serves chicken soup to Maggie because she wants Maggie to feel better and because she believes that chicken soup makes Maggie feel loved and comforted and as such promotes her health. By meeting the other two criteria, what Mom does counts as a skill. Let's call the skill: soothing a cold with chicken soup. Dr. Campbell serves Maggie chicken soup because he'd like to relieve the tenderness in her throat and he believes that the anti-inflammatory properties of the carrots and celery in chicken soup will do just this. In meeting the other two criteria for skill, Dr. Campbell has a skill, too and that skill is: soothing a cold with chicken soup. By individuating skills extensionally, despite the fact that Mom and Dr. Campbell have different reasons for acting, given that they do the same thing, they can be said to have the same skill.

There are two additional considerations that weigh heavily in favor of individuating skill instantiations in a Davidsonian manner. First, individuating skills in this way opens up the possibility of typing skills in a more intuitive and elegant way than if we did so by the reasons for which they were performed. Second, it helps to solve the mereological problem of skill possession. As we will see when we get to Criterion (3), in order for some action to qualify as a skill, that action must be learned as a result of practice. This entails that the same skill must be performed repeatedly. In order for this to be possible, we must be able to abstract away from the particular reasons one has for performing the skill at a given time. After all, if we chose to individuate skills according to concomitant reasons, then if two instantiations of a skill were not performed for the exact same reasons, then they could not qualify as the same skill. However, when we individuate skills extensionally, we can identify the following two actions as instances of instantiating the same skill: (1) I ride my bike because I want to exercise, and (2) I ride my

bike because I want to spend time outdoors. In both cases, the skill I perform is riding my bike, even though my reasons for doing so differ.

Lastly, by regarding as identical various descriptions of events that have identical extensional boundaries, it becomes possible to claim that the complex skill of riding a bike is identical to the aggregate of these more basic skills: balancing on two wheels, keeping one's feet centered on two pedals, cycling one's legs with appropriate force, keeping one's upper body still while moving one's legs, and so forth. That is, though it is possible to describe an event as one complex skill or as several basic ones, so long as the conjunction of those basic skills is extensionally identical to the description of the more complex one, then they need not be identified as different skills. Of course, we should want to know how many skills one really has and the best way to do this is to ask if each action can individually meet the above criteria for skillhood. To the extent that actions meet the criteria for skillhood and also describe extensionally identical events, they constitute the same skill. If various descriptions of events entailed the existence of different skills, then such a move would be impossible. In this way I am advocating that we are nominalists in terms of our skill-part to skill-whole relations: like a class is comprised of students, a complex skill may be a composite of sub-skills.

To sum up, individuating skills extensionally and conservatively allows us to have a coherent account of both intentional actions and skill instantiations. This choice makes it possible to give the criteria for individuation, since we avoid problems of opacity and we open up the possibility of typing skills such that having different reasons for them does not entail the existence of different skills. So, we can keep from multiplying entities since descriptions of the conjunction of various skills may be identical to descriptions of more complex skills.

3.5 Individuating skills according to their instantiations

It is important to say something about the relationship between individuating a skill and individuating an instantiation of a skill. According to my theory, an instantiation of a skill is not identical to the skill itself. This must be the case if we are to understand a skill to be (i) instantiated at different times, (ii) learned over time, and (iii) improved after its acquisition. If a skill was nothing but its instantiations, then every skill instantiation would constitute a distinct and individual skill. In this case, no two instantiations could be instantiations of the same skill. Like Heraclitus' river, one could never perform the same skill twice. But this is clearly absurd. When I read today and when I read yesterday, I am obviously performing the same skill, even though the words that I read may differ.[146] A better way to talk about my activities on these two days is to say that I am exercising the same skill today as yesterday—namely, my reading skill.

Though a skill is not identical to its instantiations, the typing of a skill may be derived from the description of the instantiation of that skill. [147] So, from the individual instance of playing a particular piece of music on the piano, we can attain the description of the type of skill that is that is being instantiated—the skill of playing the piano harmoniously.[148] This is a good move, mainly because the alternatives are worse: there is no easy way to individuate skills by way of brain processes[149] and, behaviorism and dualism are too extreme to even consider. That is, if we are not behaviorists claiming that a skill just is skillful behavior and if we are not dualists claiming that a skill is something totally different from its instantiation, then it is natural

[146] Also, if we were to individuate skills according to their instantiations, we would be faced with a slippery slope. We would be in the precarious position of defining "now." If I am reading and take a lunch break, and then return to my reading, am I instantiating one or two reading skills? What if I take a break to answer my phone? To look out the window? Do we bracket instantiations by 24 hour periods? 2 minute periods? 45 second periods, etc.?

[147] The scope is clearly an issue here, but as the success criterion will show, identification depends on the circumstances.

[148] If I only knew how to play this one particular piece of music, we may want to say that the skill is just the skill of being able to play this piece of music on the piano. As I've indicated, I think that contextual considerations will be relevant for resolving this issue.

[149] As David Foster Wallace, "Federer as Religious Experience," New York Times, August 20, 2006, Play Magazine, New York edition describes, a procedural articulation of a skill would need to account for an inordinate number of variables. This does not mean that such a description is impossible, in principle, but certainly impossible for me to present here. See Chapter five, section 1.

to type skill-kinds by way of external, skillful behavior. In fact, describing ontologically general entities according to their particular manifestations is quite common. We type universals by way of their particulars and, also, we say that a belief may be expressed behaviorally in different ways and times.[150]

But what kind of thing is a skill, if it is not its instantiation? What is its ontological status? As a first attempt to define the ontological category of skill, I would say the following: a skill is a *procedure* that is instantiated *intentionally* (criterion (1)), *successfully* (criterion (2)), and acquired through *practice* (criterion (3)). So, even though a skill is not identical to its instantiations or its developmental history, it is identified by way of them.

A separate issue that arises when a skill is not identical to its instantiation is the problem of identity over time. That is, when does a skill come into being, when is the same skill instantiated at different times, and when does it become another skill? When it comes to individuation of this sort, unfortunately, I am convinced that there is no alternative to a pragmatic criterion: the circumstances in which the skill is instantiated and identified must dictate the terms of its individuation.[151] As Hawley has argued, "a sharp criterion for what counts as the same task and what doesn't would require an unfeasibly sharp distinction between what counts as the same circumstances and what does not".[152] This means that skills exist and are typed relative to certain contexts. Though there is a fact of the matter as to what event is occurring at a specific time, the description of it—and the boundaries of its existence—cannot be identified independent of our interest. The absolute identification of a skill into principled categories and kinds is impossible. This may appear unsatisfying, but the fact of the matter is

[150] Including, of course, speech behavior.
[151] See the following section for ways of determining success.
[152] Hawley, 21.

that the identity of anything over time is highly problematic.[153] That is, the questions, "when is *p* the same person as *q*?"[154] Or "when is *o* the same object as *r?*"[155] Or "when is *k* the same piece of knowledge as *m*?" Or "when is *b* the same belief as *c*?" all lack plausible, universal answers. Even if we could identify types in an essentialist manner, as Saul Kripke has famously suggested,[156] not even he thinks that it is possible to identify when the same object of that type persists through time.[157] That is, there are only highly contestable criteria by which to determine identity over time. For this reason, I do not take it to be an especially pressing objection that we encounter this problem with skills as well.

3.6 Consciousness

Must reasons be conscious in order for them to satisfy the first criterion of skillhood? As Carl Ginet observes in relation to considerations about propositional thought, there are often propositions that I think, but that I do not contemplate in any conscious way.[158] I walk to the door to open it, I know that there is a door without having to do anything like consciously think "there is a door over there." I can want to open the door without consciously thinking, "I want to open the door." I can believe that walking over to the door, unlocking it and turning the handle is a way for me to open a door without ever, explicitly or consciously, contemplating any of these propositions. The fact is, that we act for reasons all the time without those reasons being conscious. This does not make those reasons any less reasons, nor the actions, which follow from them, any less intentional.

---

[153] See Peter Geach, "Identity," *Review of Metaphysics* 21 (1967): 3-12.
[154] See Hume, 251-63.
[155] Recall the Ship of Theseus. See Stanford Encyclopedia of Philosophy, "Material Constitution, "http://plato.stanford.edu/entries/material-constitution/," section 1, para 3.
[156] Saul Kripke, *Naming and Necessity* (Cambridge, MA: Harvard University Press, 1972).
[157] Saul Kripke, class notes, 2003.
[158] Ginet, Carl. *Knowledge, Perception and Memory* (The Netherlands: D. Reidel Publishing Company, 1975), 7.

The fact is that there is no good reason to think that a conscious thought ceases to be a thought once it ceases to be conscious. This is because the two properties in question, intentionality and consciousness, are neither identical nor interdependent. Just like there is no reason to conclude that the pretty little girl ceases to be pretty once she ceases to little, there is no reasons to think that a conscious thought ceases to be a thought, once it ceases to be conscious. Unless two properties are identical or their existence is interdependent, the absence of one does not entail the disappearance of the other.[159]

As such, in the absence of an argument to the contrary, and in the presence of so much evidence that indicates that nonconscious intentional states are plentiful, I find no reason to require that reasons must be conscious in order for them to function as reasons. Such a requirement not only mystifies the intentional behavior of animals and children, but of adult persons, too.

3.7 Propositionality

Another concern about the requirement that a skill is an action that one performs for reasons is that if reasons are beliefs and desires and we take beliefs and desires to be propositional states, then animals and small children that are incapable of propositional thought will thereby be incapable of skill possession. I think that the best way to respond to this concern is to note that intentional states need not be cashed out in terms of concepts and propositions. As I elucidated in chapter two, section four, there are several options for exploring the nature of intentional content that do not construe cognitive states as necessarily constituted by concepts.

[159] And in fact, those arguments must respond to the abundance of evidence that exists showing a dissociation between intentional states and consciousness. That is, even if conscious states are always intentional, this does not mean that the two are identical. For there can still be plenty of intentional states that are not conscious. See David M. Rosenthal, "Being Conscious of Ourselves," *The Monist* 87 (2004): 162; David M. Rosenthal, "State Consciousness and Transitive Consciousness," *Consciousness and Cognition* 2 (1994): 356-8.

For example, we can appeal to a P.F. Strawson-type feature-placing language in order to account for the pro attitude aspect of reasons. In this way, a desire can be for "that," where "that" is not a conceptualized, identifiable or reidentifiable particular. Instead a pro attitude can be for an embedded feature of a perceptual array. So, we can ascribe the desire for those berries over there to a creature, without requiring that the creature possesses a general and recombinatorial concept of "berries," as such.

Further, when considering the instrumental belief portion of a reason for action, we can appeal to an affordance based instrumental belief like the kind advocated by Michael Dummett and Ruth Millikan. We can also think of partially recombinatorial, domain specific, instrumental beliefs such as those described by Susan Hurley. Or we can appeal to a José Bermúdez type proto-logic where causality and contrariness constitute a basic kind of conditional reasoning, and as such, can offer a nonconceptual theory of instrumental belief.

Any of these accounts can be worked out to provide a nonpropositional theory of reasons for action. As such, we can be sure that simply requiring that a skill is an intentional action does not entail that only adult humans can be in possession of them. By looking to the proto-level of desire and belief, we can construct reasons for actions that are not fully abstract, general or recombinatorial in nature.

4. Criterion 2: ability

Criterion 2) If *a* is a skill of *S*'s, then under conditions *c*, *S* is usually able to instantiate *a* successfully.

In this section, I would like to turn my attention to abilities. As I argued above, skill is a species of ability, though not all abilities are skills.[160] This means that in order for *a* to qualify as

[160] Recall my lawn mower: it is able to cut grass, but it does not know how to cut grass; it does not possess grass-cutting skills. See also Carr, "The Logic of Knowing How and Ability," 409: "Reports of knowing how differ from

a skill of *S*'s, *S* must have the ability to instantiate *a* successfully. We can refer to Criterion (2) as the ability or success condition for skill. In examining the commitments and limitations of Criterion (2), I will focus on the following issues:

(1) Skill is a success term.
(2) Success attributions assume that certain normal underlying conditions obtain; I will call these opportunity conditions.
(3) The standards of success for an ability are context-dependent.
(4) Success is not required 100% of the time.
(5) Treating abilities in straightforward counterfactual terms allows us to avoid some of the hairier issues associated with dispositional states and events.

4.1 Success and ability

It has sometimes been suggested that knowing-how cannot require ability since there are cases where knowing-how attributions are obviously warranted, but a subject is unable to perform the ability successfully. Carr states,

> In the first place, although it might seem reasonable to argue that possession of physical ability must be considered a necessary condition of knowing how to do something, it is clear, surely, that there is nothing in the least paradoxical about describing an elderly and arthritic piano teacher or a temporarily incapacitated gymnast as knowing how to do whatever they cannot currently perform…present possession of the ability does not appear to be necessary.[161]

Along with the arthritic piano teacher and the incapacitated gymnast, I'd like to say that the amputee cyclist and me, when I'm in a dark room, still have the respective abilities of playing the piano, doing gymnastics, riding a bike and reading. In these cases, we all know how to *a* without being able to *a* successfully, right now.

But does the inability to perform a skill now indicate that success is not a necessary criterion for skill? I will argue that the answer to this question is "no." For, despite the fact that

---

those of ability in taking members of a different class of action descriptions as their objects. The descriptions in question are those of intentional actions rather than mere instances of agent causation."

[161] David Carr, "Knowledge in Practice." *American Philosophical Quarterly* 8 (1981): 53

counterexamples seem to abound, the ability to successfully perform *a is* a prerequisite for knowing how to *a*. In order to reconcile this apparently contradictory position, I will appeal to Katherine Hawley's account of knowing-how, where she claims that "actual success is not a necessary condition for knowledge-how, but success under certain counterfactual circumstances is indeed necessary."[162] What makes ability attributions true is not that one can perform successfully under *any* circumstances, but that one can perform successfully under circumstances that are presumed to be *normal.* When circumstances are not normal, through the use of counterfactuals, we can assess whether a person would, under normal circumstances, be able to perform the ability successfully.

The above examples help us to see that the reason failure doesn't count against a subject's possessing a particular skill is because the conditions in which they are executed are abnormal. We say that the arthritic piano teacher has the skill of playing the piano because we consider normal conditions for playing the piano, having fingers that can move with dexterity. If the piano teacher satisfied those conditions, then indeed he'd be able to play the piano successfully. In the same way, we see that the amputee cyclist knows how to ride a bike, because if he had two legs, he'd be able to ride a bike successfully. [163] And also, I know how to read, even in the dark, because if it were light, I'd be able to read. After all, no one believes that to know how to read requires the ability to read in all circumstances. If I don't have my reading glasses, or the proper lighting, or access to reading material, I cannot read, but I still possess the skill of reading.

---

[162] Hawley, 20.

[163] "[A]fter losing his leg, Seth does not know how to ride a bike under the circumstances of having one leg. But neither did he know how to ride a bike under the circumstances of having one leg *before* he lost his leg (Williams, section 7).

For every ability that may be successfully instantiated, there is a set of background opportunity conditions that we assume must be satisfied in order for failure to count against skill possession. If these conditions are not met, then failure to *a*, does not preclude the skill of *a*-ing; rather we should understand it as lacking the opportunity to *a*. As Hawley states, "failure to perform in abnormal circumstances for that task doesn't usually count against someone's possessing knowledge-how."[164]

For example: Knowing how to ski involves the opportunity conditions of being on a mountain covered with snow, having skis, having two legs, having a basic level of coordination, etc. If the opportunity conditions are not satisfied, then not being able to ski successfully, will not count against having the ability to ski, just so long as if those conditions were met, then one *would* be able to ski successfully.

Recall Mary Lou and Bela.[165] Bela has oodles of propositional knowledge, but does not possess skill and Mary Lou has less propositional knowledge, but possesses the skill that Bela lacks. Recall, however, that neither Bela nor Mary Lou can presently perform a standing layout on beam. We can now explain why we attribute skill to Mary Lou but not to Bela. Mary Lou has skill because she *would* be able to successfully perform a standing layout on beam if the conditions were normal. That is, if Mary Lou met the opportunity conditions of weight, strength and flexibility, then she'd be able to perform the skill successfully. On the other hand, even normal circumstances for performing a standing layout on beam cannot provide Bela with the ability to perform successfully.

By appealing to success under counterfactual conditions, I can do justice both to the idea that ability is a success term and explain why success does not always require success now,

[164] Hawley, 22.
[165] See Chapter three, section 3.

under any and all conditions. By following Hawley and requiring that knowing-how requires success under normal counterfactual conditions, it is possible to retain this analytic aspect of ability. With this approach, we can explain why the ability to *a* successfully is required for possessing the skill of *a*-ing, while also accounting for why apparent counterexamples are just that—apparent. Of course, "normal" conditions are notoriously difficult to specify and though it would be convenient to avoid this kind of talk, I do not think that it is at all possible. The closest that we can get to an accurate account of ability will involve the messy and complicated conditions of the world.

Problematically, it may seem that we can present the counterfactual or opportunity conditions in such a way so that everyone knows how to do everything.[166] So, we can say that I know how to fly a plane, because if I went to flight school and practiced flying for 80 plus hours, etc., then I'd be able to fly a plane successfully. However, it would be rather odd to build the learning conditions for a skill into the opportunity conditions for possessing it. As such, I think we have a natural limitation on the types of counterfactual circumstances that we ought to bring to bear in determining normal counterfactual circumstances. Quite clearly, any sort of learning, training, drilling or developing conditions cannot be legitimate posits for assessing counterfactual success. After all, we want to determine which skills agents really possess, not the skills that they might possess, if they were to learn them.

4.2 Standards for success

It is important distinguish between (i) the opportunity conditions that we consider normal for performing *a* successfully, and (ii) the standards by which we judge that *a* has been instantiated successfully. A discussion of opportunity conditions focuses on existent underlying conditions that must be satisfied in order for *S's* failure to *a* to count against *S's* ability to *a.*

[166] See Hawley, 23.

Opportunity conditions are the background conditions that we consider normal for *a-ing*. Standards for success are just that: standards for judging when *S a*'s successfully.

To my disappointment, it turns out that the standards for success are no more easily or neatly identified than are "normal conditions." This means that, unfortunately, for most skills, we should not hope to find universal standards by which we can determine successful performance. As Hawley states,

> For many families of tasks, it will be difficult to identify a plausible candidate for canonical circumstances or standards…in a UK context, it would be reasonable to infer from Sarah's knowing how to drive that she knows how to drive a manual, stick shift car. In most US contexts, however, this would not be a reasonable inference. To take a different example, a child might be said to know how to cook if she knows how to use the stove safely, whilst we would set standards higher…for an ordinary adult's "knowing how to cook," and set them higher still when discussing a chef in training.[167]

Success attributions depend essentially both on the context in which one attempts to *a* and on the agent who is doing the *a*-ing. For instance, we have different standards of success for agents who occupy different age groups, genders, and levels of expertise. Likewise, standards of success vary according to geographical and temporal location, season, altitude, etc. We see this variability in standards represented quite straightforwardly in the structure of athletic competitions. Athletes compete only against other athletes who are in their own skill level: in karate, brown belts do not compete against yellow belts and in gymnastics, level I gymnasts do not compete against level II athletes. High school students do not compete against college athletes and professionals do not compete against novices.

We expect experts to perform at higher levels than enthusiasts. If an elite runner runs a mile in over 4 minutes, we call his performance a failure; but if an enthusiast runner performs this very same feat, we consider his performance to be a great success. We should not expect

[167] Hawley, 21-2.

that when a child is successful at playing the piano that she is able to play to the same standard as Denis Matsuev. But she can play the piano nonetheless—she strikes the correct keys, reads notes, keeps time, etc. The reason why a child can play the piano successfully, even when playing poorly compared to a professional pianists, is that the standards of success for the child and for the professional differ. It would be rather strange indeed, if in order to play the piano, a child would have to play Chopin's Nocturne Op. 9 no.2 flawlessly and it would be even stranger, perhaps, if Matsuev could be said to successfully instantiate his piano-playing skill after playing "Mary had a Little Lamb." The reason that neither of these judgments seem right is because they assume that the standards of success for one agent must be identical across all agents when, in fact, those standards ought to be sensitive to context.

Importantly, though the standards of success for most skills are not context-independent, given particular conditions, standards that are appropriate relative to a context, quite naturally arise. So, though standards for success are not absolute, they are not absolutely relative, either. Once a context is determined, it becomes relatively unproblematic to identify the relevant standards of success. After all, isn't that what teachers do every time they give an exam or grade an assignment? And simply because the standards of success for a Freshman undergraduate class are not the same as those for a graduate seminar, this does not entail, of course, that neither class has any standards for success, nor that those standards are completely subjective.

4.3 Reliable Success

We do not want to say that in order to have the ability to *a*, *S* must *a* successfully 100% of the time. As Hawley states, "Knowledge-how requires reliable, but not exceptionless

success."[168] This fact is reflected in Criterion (2) with the word "usually." Instead of requiring unqualified success, I claim that ability requires that *S* is *reliably* able to *a*.

After all, even if I trip once in a while, this does not mean that I do not know how to walk, and simply because Joe gets into a car accident does not mean that he does not know how to drive. Of course, if Joe frequently gets into car accidents, then we should want to say that he does not possess the skill of driving. But where do we draw a line? Unfortunately, as with "normal conditions" and "standards of success," "reliably" cannot be pinned down in anything like absolute terms. We cannot say that if I trip six times a week then I possess the skill of walking, but if I trip seven, then I don't know how to walk. Again, we must step into the murky waters of this messy world. But this fact should not make us too worried. After all, there are quite clearly men with hair and men that are bald. Simply because we cannot solve the Sorites paradox[169] and identify the exact number of hairs that makes a man bald rather than not bald, does not mean that we cannot meaningfully and accurately use these terms to refer to men with particular properties. Likewise, just because we cannot pin a number onto "reliably" for a particular skill, does not mean that we have no way of differentiating when an ability is performed this way rather than another. So, though as philosophers we should want necessary and sufficient conditions for when one reliably drives or walks or bakes, we shall have to get by with the best we have, which is to judge according to context and common sense. To stay on the safe side, however, all of the hypothetical examples that I will make use of will not fall into the ambiguous area of reliable v. unreliable success. I will stick to examples that clearly do or do not meet this reliability condition.

4.4 Counterfactual v. Dispositions

[168] Hawley, 24.
[169] See Stanford Encylopedia of Philosophy, "Sorites Paradox," http://plato.stanford.edu/entries/sorites-paradox/, sections 1 & 2.

The last thing to notice about criterion (2) is that it treats ability as success under counterfactual conditions, not as a disposition to succeed. There are a couple of subtle advantages to this approach: (1) We do not need to speak about categorical base structures, (2) we do not have to claim that certain conditions elicit or trigger certain kinds of behaviors, and (3) counterfactuals seem to align nicely with the way people naturally think of skills.

It is often considered imperative that an explanation in terms of dispositions makes reference to some actually present, structural base property of the system to which they belong.[170] That is, it is thought that dispositions are true because of the very real and present properties of an object or system now; they are true in virtue of their being grounded in the structural properties of that system. For example, we should want to know what the property of a lamp is now such that if it is dropped, it will break. We should want the dispositonal state of being fragile to have an underlying explanation in terms of the basic structure of porcelain, not only at the time of breaking, but at the time before breaking, too. That is, at the time that a lamp is standing on a credenza and one is warning her five-year-old niece to be careful because the lamp is fragile, at that time, we should want to know what it is about the lamp that makes it the case that it would break if it were dropped. Fortunately, with the property of being fragile, such an explanation is forthcoming: it has to do with the loose spatial arrangement of molecules. But such an explanation is much trickier when it comes to abilities.

After all, it seems perfectly reasonable to hold that there is no one particular, easily identifiable, occurrent property of my brain, right now, that makes it the case that I could succeed in riding a bike tomorrow, if I tried. Even though when I ride a bike this ability is most certainly due to a whole host of facts about my brain and body now and those facts are

[170] See Stanford Encyclopedia of Philosophy, "Dispositions," http://plato.stanford.edu/entries/dispositions/, section 3.

determined by my body chemistry, brain structure, and past experiences, along with a whole array of other conditions and facts, it does not seem to be necessary or even plausible to attribute to me today, anything like a structural property in virtue of which my ability to succeed in riding a bike tomorrow is grounded. It seems unnecessary to say that there is anything stored or structured or occurring in me now, beyond the potential for some procedure to be instantiated tomorrow, that explains, in the way that the spatial layout of molecules does for the disposition of being fragile, my ability to ride a bike.[171]

But, just because it seems unlikely that we will find any identifiable, ostensible property to ground the counterfactual about my ability to ride a bike, that does not make my ability magic. After all, it is either true or false right now that either it will rain or it will not rain on November 19, 2020. And this fact is either true or false due to the state of the world now along with a whole host of other natural facts. But simply because there are facts about the world that make this counterfactual true or false does not make it the case that this counterfactual is a disposition. While everything that will be true in the future will be true because of the present state of the world, this does not entail that everything that will be true in the future will be true solely in virtue of the dispositions of objects and systems.

Quite simply, the counterfactual route does not require us to speculate about the nature of the facts today that make it the case that I'll be able to ride a bike tomorrow. Rather, with a counterfactual analysis we can stay neutral on the nature of the procedures responsible for successful ability instantiation when those procedures are not being instantiated. This is an advantage of counterfactuals since the facts that follow from both dispositions and counterfactuals are true of the objects and systems to which they are attributed before they are

[171] See chapter five, section 1. Procedures themselves are complicated enough to try to identify, never mind trying to identify the facts that make it the case that these procedures will be instantiated before they actually are.

manifest. It is true of me now that if I try to ride my bike tomorrow, under normal conditions, I will succeed. Counterfactuals afford the same central prediction workload of a dispositional account, but also avoid substantive commitments on issues that we simply cannot judge.

Further, by sticking to a counterfactual analysis, we can avoid the embarrassing problem of having to identify normal circumstances with triggering conditions. With dispositions, we want to say that particular conditions are not just required, but responsible for the expression of a disposition. For instance, sugar dissolves when placed in water, because it is soluble. Placing sugar in water is not just necessary, but sufficient for the actualization of the dispositional property; the conditions *explain* why the sugar dissolves. But the same cannot be said of the normal background conditions necessary for the successful instantiation of ability. After all, it would be rather strange to say that the presence of a balance beam elicits Mary Lou's ability to do a standing layout; she could not be said to instantiate her ability *because* there was a balance beam in her presence. As Wallis states,

> [A] dispositional analysis holds that one is disposed to behave in a certain way in the specified conditions. Task-specificity (a success under counterfactual conditions type analysis), in contrast, states only that certain conditions underlie the manifestation of knowledge-how, regardless of whether they are sufficient to elicit its manifestation. So, task-specificity asserts that Bill requires access to conventional oven in order to manifest his cake-baking know-how. It does not claim that access to a conventional oven will elicit Bill's cake –baking behavior (parenthetical not in original).[172]

While we could say that the only problem here is describing all of the normal conditions in enough detail to ensure their sufficiency for triggering an ability, it seems that this route just clings to a dispositional account for sentimental reasons and not because of theoretical superiority. After all, Ryle himself admits that it is impossible to determine when, in fact, a subject possesses an ability, because we can never identify with certainty the conditions that

[172] Charles Wallis, "Consciousness, Context and Know-how." *Synthese* 160 (2008): 144.

would be sufficient to elicit the performance of that ability. The reason Ryle thinks we can never be certain of ability possession is precisely because he thinks of abilities as dispositions and dispositions must have sufficient triggering conditions.[173]

There is no reason to believe that we must have triggering conditions in order to have a full account of ability. The idea that we need an explanation of abilities that includes anything like sufficient conditions for their instantiation seems to miss the mark altogether. The question about whether one has a skill and the question of what will elicit one to perform that skill are conceptually distinct. The fact that dispositions, classically understood, blur this distinction is simply one more reason to avoid them.

Lastly, it seems that the common sense reasons people give for ability attributions in the absence of actual ability are often counterfactual in nature. If *S* cannot *a* now, and we think that *S* still has the skill of *a*-ing, many people will say that the reason that *S* has this skill is because *S* could *a* under different circumstances. As Wallis states, "if one asks people why they suppose that the blind woman can read, they invariably link the ascription to a counterfactual ability. They say, "she could read, if she had the ability to see."[174] Of course, people may say this because they believe *S* has a particular disposition to *a*, but it seems that this additional substantive posit does no real work for us. Obviously, if we can avoid complications and controversial commitments, then we ought to.

[173] See Ryle, 44.
[174] Wallis, 149.

CHAPTER 5

Criterion 3: If *a* is a skill of *S*'s at *t*, then there was some time *t*-1 that *S* could not instantiate *a* successfully, and *S* has learned to instantiate *a* successfully as a result of practice.

1. Practical Intelligence

Criterion (3) is the learning criterion for skill. This criterion ensures that skill is not just an intentional ability, but an intelligent ability. This means that the type of development that occurs as a result of practice is not simply the refinement of intentions or the accumulation of factual knowledge, but an improvement in the practical instantiation of an action. That is, intelligence guides skill not just as an impetus, but as a characteristic operative throughout its manifestation. I argue the intelligence of skill is practical intelligence; it is the sensitivity and control that characterize the complex motor patterns that are expressed in skillful action.[175]

In chapter three, I claim that knowing-how cannot be reduced to factual or propositional knowledge. I claim that knowledge of facts cannot account for the manner in which a skill is instantiated. In this chapter, it should become clear that the particular cognitive aspect of skill that cannot be accounted for by knowledge of facts or rules is practical intelligence. That is, propositional knowledge cannot account for the sensitivity and control that characterize the operations of skillful action; no amount of factual knowledge can explain the particular manner in which a skill is instantiated. Such knowledge cannot account for the precision, grace, elegance, force, etc., with which a skill is performed.

[175] I will not define cognition or intelligence here, though I wish I could. Instead, I will focus on skills that seem unobjectionably cognitive. To extend these considerations to more questionable activities, a full account of cognition would be necessary. Unfortunately, such an account is beyond the scope of this dissertation.

It is a distinct but equally untenable option to reduce the practical intelligence of skill to the intelligence of the reasons or intentions for which the skill is performed. This should be quite clear since if intentions are responsible for the particular manner in which complex motor patterns are instantiated, then it should be impossible to have differences in the manner of instantiation without corresponding differences in the intentions that instigate them. But unless we mean something quite different by "intention" than a reason or a cause for action, then this position is ineffective.

The fact is that actions can succeed or fail, but they can also be more or less precise, graceful, sharp, defined, focused, expressive, clever, elegant, fluid, appropriate, and so on. It is exactly the style or manner in which a skill is performed that is indicative of the role of practical intelligence in action. This embodied, practical kind of intelligence is paradigmatically represented in the *way* that actions are instantiated: it is displayed in the expressiveness of a piano player's performance, the power of a swimmer's stroke, the fluidity of a dancer's movements, the precision of a gymnast's landings and the spin off a tennis player's backhand. We can easily appreciate the various degrees of intelligence inherent in skill instantiation by noticing the differences in the manner in which different persons at different levels of skill perform them. As Gilbert Ryle has noted,

> The boxer, the surgeon, the poet and the salesman apply their special criteria in the performance of their special tasks, for they are trying to get things right; and they are appraised as clever, skilful, inspired or shrewd not for the ways in which they consider, if they consider at all, prescriptions for conducting their special performances, but for the ways in which they conduct those performances themselves. Whether or not the boxer plans his maneuvers before executing them, his cleverness at boxing is decided in the light of how he fights. If he is a Hamlet of the ring, he will be condemned as an inferior fighter, though perhaps a brilliant theorist or critic. Cleverness at fighting is exhibited in giving and parrying of blows, not in the acceptance or rejection of propositions about blows, just as ability at reasoning is exhibited in the construction of valid arguments and the detection of fallacies, not in the avowal of logicians' formulae. Nor does the

> surgeons' skill function in his tongue uttering medical truths but only in his hands making the correct movements.[176]

The locus of intelligent action is action and not thought; it is expressed in the manner an action is performed and the sensitivity and control that is required for the performance. Such intelligence is highly nuanced, sophisticated and complex. We can see this because very small adjustments and variations in spatial positions of bodies and objects, as well as timing can make very big differences in the manner in which a skill is performed. And it is exactly these nuances in instantiating a skill that require the sensitivity and control of embodied practical intelligence, rather than a stockpile of rules and facts. This does not mean that the precision of action is unintelligent or instinctual, but rather that it simply cannot be the deliberate process of reasoning from premise to conclusion. It cannot be propositional thought. As David Foster Wallace writes,

> By way of illustration, let's slow things way down. Imagine that you, a tennis player, are standing just behind your deuce corner's baseline. A ball is served to your forehand — you pivot (or rotate) so that your side is to the ball's incoming path and start to take your racket back for the forehand return. Keep visualizing up to where you're about halfway into the stroke's forward motion; the incoming ball is now just off your front hip, maybe six inches from point of impact. Consider some of the variables involved here. On the vertical plane, angling your racket face just a couple degrees forward or back will create topspin or slice, respectively; keeping it perpendicular will produce a flat, spinless drive. Horizontally, adjusting the racket face ever so slightly to the left or right, and hitting the ball maybe a millisecond early or late, will result in a cross-court versus down-the-line return. Further slight changes in the curves of your groundstroke's motion and follow-through will help determine how high your return passes over the net, which, together with the speed at which you're swinging (along with certain characteristics of the spin you impart), will affect how deep or shallow in the opponent's court your return lands, how high it bounces, etc. These are just the broadest distinctions, of course — like, there's heavy topspin vs. light topspin, or sharply cross-court vs. only slightly cross-court, etc. There are also the issues of how close you're allowing the ball to get to your body, what grip you're using, the extent to which your knees are bent and/or weight's moving forward, and whether you're able simultaneously to watch the ball and to see what your opponent's doing after he serves. These all matter, too.

[176] Gilbert Ryle, *The Concept of Mind* (Chicago: The University of Chicago Press, 1994), 48-9.

> Plus there's the fact that you're not putting a static object into motion here but rather reversing the flight and (to a varying extent) spin of a projectile coming toward you — coming, in the case of pro tennis, at speeds that make conscious thought impossible. Mario Ancic's first serve, for instance, often comes in around 130 m.p.h. Since it's 78 feet from Ancic's baseline to yours, that means it takes 0.41 seconds for his serve to reach you. This is less than the time it takes to blink quickly, twice. [177]

Wallace's comments elucidate our dilemma: high level tennis requires execution of skill at such a pace so as to circumvent conceptual, propositional, conscious reasoning. But these skills are so precise, so deliberate that we are compelled to attribute intelligence to their execution. Temporally, it would seem that elite athletes are functioning in the operative range of reflexes, performing purely physical reactions that bypass thought. And yet, an effective return of serve depends on an enormous set of decisions and physical adjustments that are infinitely more complex and intentional than "blinking or jumping when startled."[178]

The fact remains that embodied skills require a high level of awareness, attention and control. The slightest adjustment has enormous impacts and yet, the shear quickness of performance means that critical, immediate decisions about angle, force, timing, speed, etc., cannot consist in doing practical syllogism in one's head. It is not only that there isn't enough time, but there is no way to connect the reasoning to the implementation. That is, propositional reasoning does not have the fineness of grain to distinguish between the micro-second, micro-millimeter adjustments that end up making all the difference in the way a skill is performed.[179] But even if fineness of grain could be accommodated conceptually, as I argued in the concluding section of chapter two, the inherent particularity of these kinds of adjustments and responses

---

[177]David Foster Wallace, "Federer as Religious Experience," New York Times, August 20, 2006, Play Magazine, New York edition, para 23.
[178] Wallace, para 24.
[179] For arguments in favor of nonconceptual content based on fineness of grain, see Gareth Evans, *The Varieties of Reference* (Oxford: Oxford University Press, 1982), 229; Christopher Peacocke, *A Study of Concepts* (Cambridge MA: MIT Press, 1992); Michael Tye, *Ten Problems of Consciousness* (Cambridge MA: MIT Press, 1995); Michael Tye, "Nonconceptual Content, Richness, and Fineness of Grain," in *Perceptual Experience*, ed. Tamar Szabo Gendler and John Hawthorne, 504-530 (Oxford: Oxford University Press, 2006).

precludes them from following from practical syllogism.[180]

To clarify, by "practical intelligence" I do not mean Aristotelian practical reason.[181] Whereas practical reason is reasoning that is directed at the realm of action, practical intelligence is intelligence that characterizes the manner in which an action is performed. That is, practical reason is *about* action, but practical intelligence *governs the instantiation of* action. So, while practical reason can tell us if and when an action ought to be performed, practical intelligence determines *how* it is performed—with style or grace, sloppily or sporadically.[182]

The way in which an action is instantiated, the sensitivity and control that is operative in its success, must be attributed to the psychological capacities of an agent. Skill is clearly more than a matter of brute instinct or mechanistic function. But there is no reason to posit an additional set of intentional states, which are behind, beyond, or above the procedural motor patterns that are expressed when a skill is instantiated. In fact, it is the central point of my non-reducibility argument of chapter three that such a posit cannot account for the practical intelligence of skill.

I propose that practical intelligence, or the sensitivity and control that is learned through practice, is a feature of skill procedures. As such, the cognitive aspect of a skill is not a conjunction of a thought and the operation it governs, but rather, skill is an operation of a particular kind; namely one that is characterized by practical intelligence. Practical intelligence, then, is best understood as procedural knowledge. It is more likely to be formalized in a function or an algorithm, than in propositional form.

---

[180] See chapter three, section 4.1.
[181] Aristotle, *The Nicomachean Ethics* (Oxford: Oxford University Press, 1988), 1140a25-1140b30.
[182] As Ryle, 32 has argued, the intelligence of practice is not the intelligence of thinking propositions and also doing something, but rather, doing something in a particular way, he states that to be intelligent "my performance has a special procedure or manner, not special antecedents…It is to perform one operation in a certain manner or with a certain style or procedure, and the description of this *modus operandi* has to be in terms of such semi-dispositional, semi-episodic epithets as 'alert', 'careful', 'critical', 'ingenious', 'logical' etc."

To be clear, my definition entails that a successful, intentional action, or any action that satisfies the first two criteria, but not the third, is not a skill. This is because it is possible for an action to be both intentional and successful and yet lack the deliberateness and control that is characteristic of skill. Recall my making a basket by throwing a basketball over my head. Surely my action is intentional and successful (I tried to make the basket and I did make it), but it is also just lucky. I do not exhibit a behind-the-head-basketball-shot-skill because the manner in which I perform the action is not guided by practical intelligence; that is, it lacks the kind of sensitivity and control of attention and response that is characteristic of intelligent action—and that I claim is developed through practice.

2. Defining practice

To begin, I'd like to present a basic working definition of practice. In very general terms, I claim that practice requires repeating an action for the sake of learning or improvement. So, we can stipulate that practice results in a skill when repeating an action for the sake of improving the likelihood that the action will be performed successfully, successfully improves performance.

The requirement for repetition as a feature of practice should be fairly unproblematic, for we wouldn't call an activity practice if it were done only once. The idea that practice is performed with the intention to improve should be fairly straightforward as well. It is this feature of practice that separates it from mere repetition. After all, even if a heart improved at pumping blood because of repeated pumping, we would not say that the heart was practicing pumping blood.

However, to be forthright, the above options, are not exclusive. One may wish to define practice as the repetition of an intentional action, but not require that the intention for which the action is performed is an intention to learn or to improve.[183]

I advance a somewhat stringent standard for practice in order to avoid unnecessary distraction. This is because, in the absence of a definition of cognition, it is wise, I think, to keep the kinds of activities that I am dealing with fairly unobjectionable. I claim that in order for some activity to count as practice, that activity must be performed with the intention to learn, to figure something out, or to improve. This keeps the discussion rooted in noncontroversial claims about the limits of practice and the cognitive activities involved therein. Whatever else may qualify as practice after we flesh out a whole host of assumptions, implications and commitments about intelligence, cognition and learning, the kind of repetition that is performed for the sake of improving describes robustly intelligent processes and, thus, produces potentially cognitive results.

Though in the broader scheme of things, it will be necessary to decide which activities count as practice, as opposed to mere repetition, so that we can determine if their results may be related to perception in such a way as to constitute an instance of cognitive penetration, I need not answer this question here. In order to sketch the relationship between skill and perception that must obtain in order for it to qualify as an instance of cognitive penetration, I need not decide the number of things that may potentially enter into that relationship. This is a related, but distinct question.

To return to the intentional requirement of practice, let me emphasize the following point: though it is required that repetition is undertaken with an intention to learn or improve, the

[183] I will discuss the implications of this alternative way of framing the notion of practice in my concluding chapter six, section 1.1.

intention to learn need not necessarily correspond with the procedure learned. So, repeating some action with the intention to improve at *x* does not necessitate that it is *x* that is learned as a result of the practice. My claim is that the particular procedure learned through practice need not correspond to the intention for which the practice is performed, just so long as an intention for learning is present and the learning results from the repetition of the action that is performed with that intention. So, an action that is performed for the sake of learning *x* may result in the learning of *y* and just so long as the intention to learn is present, the fact that another thing is learned becomes irrelevant. Though, it makes sense to think of *y* as somehow related to *x*, and as such, the intention for *x* may somehow tacitly include the intention for *y*, I do not think that this must be the case. When *y* is learned as a result of practicing *a*, when this practice is caused by the intention to learn *x*, where *x* is not equal to *y* and not obviously related to *y*, we can still hold that *y* is learned as a result of practice.

To elucidate, take the following example: I may intentionally jump rope every Thursday and do so with the intention of improving my ability at jumping rope. As a result of this practice, I may learn how to hop on one foot in a straight line. Though my intentionally jumping rope was not for the explicit reason of learning how to hop on one foot in a straight line, because the learning was intentional and repeated, and resulted in the improvement of the probability of my successfully hopping on one foot in a straight line, then it is appropriate to say that I learned this skill through practice.

Very generally, practice can be thought of as a process of practical trial and error. By acting with the intention to learn, one engages in a process of figuring out what works for oneself; one engages in a process of practical problem solving. This process can be described, not at all metaphorically, as an agent attempting various strategies in order to exploit those

methods, efforts and variations that prove successful. When one succeeds, one repeats the things that she has done in order to achieve that success. When one fails, one avoids repeating those things that lead to failure. In the process of developing practical intelligence, an agent figures out which features of an action-space to attend to, how much force should be exerted at which stages of performance and which expectations are appropriate given the stimuli, etc., by deliberately attending and responding in different ways to various inputs and outputs. In practice, one tests the various degrees of attention and exertion that are necessary in order to successfully instantiate a skill and often, one repeats what works until it becomes second nature.

3. Practice and practical intelligence

Now that we have a working definition for practice, I'd like to return to the relationship between practice and practical intelligence. Let us begin all the way at the beginning by asking, "Why practice?" The most obvious answer is, "to learn" or "to improve." But then we should ask, "To learn what?" and here is where the connection between practice and practical intelligence becomes apparent: through practice, a subject learns how to effectively instantiate complex motor patterns. A subject learns to be attentive and responsive to the right things, at the right times, in the right ways. After all, this is what acquiring a new skill requires.

My claim here is that the way a skill is acquired is indicative of what is learned as a result of that particular method of acquisition. So, the types of methods that are used to learn a skill can help us to understand the nature of the things that are learned as a result of those methods.

After all, knowledge can be acquired in a multiplicity of ways. One can learn what the temperature is outside by being told; one can learn what the color of a new car is by looking; one can learn the capitals of US states by memorizing; and one can learn to swim by practicing. The

appropriateness of a particular method of learning is instructive in elucidating the nature of that which is learned as its result.

Now, since doing is imperative to skill learning, we should be aware that we are dealing with the realm of action and not just in the realm of thought. This feature of skill acquisition should alert us to the fact that what is acquired in practice is not only the knowledge of when and in what way to respond in order to be successful, but the actual development of appropriate control, timing, force, precision, etc., necessary to instantiate an action. This practical knowledge cannot be accounted for with intentions or facts.

After all, intention alone cannot guarantee success; reasons for performing an action do not comprise the capacity to perform that action successfully. In order to do this, one must actually develop the capacity to perform the action appropriately: that is, one must practice. If a skill, merely through its intention, could be performed with the appropriate attention and control, then practicing that skill would be redundant. Additionally, it cannot be that intention itself is learned as a result of practice, since practicing a skill presupposes that the intention to perform the skill already exists. An explanation of the effectiveness of practice that appealed to the refinement of intentions would be patently circular.[184]

Further, practical intelligence is obviously quite different from learning a fact or a rule about a skill. For instance, it is different than learning the fact that one should be responsive to *f*. If practical intelligence were nothing more than this, then the appropriate manner of learning would not be practicing the skill, but memorizing the instructions governing it. Practice, however, requires a type of drilling, training and repetition different from fact-oriented repetition: often, when an embodied skill is practiced, it is repeated until it becomes automatic; it is drilled until it becomes second nature. There does not seem to be a parallel between this sort

[184] Or we would have to mean something very different by "intention" than we usually mean.

of drilling and the learning of facts or rules. Though memorizing a fact may require repetition, one need not, like in skill, repeat the fact after one has learned it.[185] The style of the fact that is memorized does not change with repetition.

Also, the way in which one engages in practice strongly suggests that what is learned as a result of practice is cognitive or intelligent. That is, when one performs an action with the intention to learn or to improve, then one is engaging in a deliberate, practical trial and error process that involves, willing, feeling, understanding, attending, drawing connections, discriminating, controlling, distinguishing, etc. Learning through practice not only involves a whole host of sophisticated cognitive functions operative during the learning but also results in an increase in the sensitivity and control that is characteristic of the operations constitutive of skillful action. As such, learning through practice entails that that which is learned is not propositional, since if it were then practicing it would be irrelevant for acquisition, but it also ensures that that which is learned is cognitive, since it is learned as a result of paradigmatically cognitive processes.

Given that learning a skill through practice involves a high level of cognitive engagement and sophistication, it should be clear that the learned content is legitimately intelligent. After all, if we substitute beliefs for actions, the method of acquiring beliefs via an intentional learning process seems to be the most clearly intelligent way in which to form them. This should make us confident that the results of intentional learning, whether intentional states or processes, are squarely cognitive.

One could, of course, claim that since the results of practice are not intentional states like beliefs and desires, but rather, abilities to instantiate complex procedures under normal

[185] See chapter 3, section 3.2 for a discussion of how the necessity of being in "the zone" for high-level athletic performance has implications for the nonpropositionality of knowing-how.

counterfactual conditions, then these are not the right kinds of things to be characterized by intelligence in the first place. One may claim that I am committing a category error; like claiming that Tuesday is yellow. In that case, my argument from analogy to beliefs would not work because the only kinds of things that could be legitimately cognitive would be intentional states. However, this sort of belief/desire chauvinism leads to some thoroughly unacceptable consequences. The fact is that either we accept that nonintentional states or events may be intelligent, or we relegate all processes from tennis-playing to problem-solving to the realm of the noncognitive.[186] In addition, we relegate all the behaviors of non-linguistic animals and pre-linguistic children to the world of brute, mechanistic force.[187]

This follows because there is no non-circular way of attributing beliefs and desires to animals and small children. If intentional states are required for intelligence, then in order for animals and small children to count as intelligent, they would have to possess intentional states. But the only way we have to attribute intentional states to animals and pre-linguistic children, is by first establishing that their behaviors are sophisticated enough to warrant the attribution of beliefs and desires. As such, the cognitive quality of a behavior is identified as a prerequisite for attributing beliefs and desires to non or pre-linguistic creatures, but it is these very beliefs and desires that are supposed to establish whether the behavior is intelligent or not.

By claiming that only beliefs and desires can be characterized by intelligence we not only (1) beg the question against the possibility of intelligent processes (What would make a process intelligent? Nothing: processes aren't the kinds of things that can be intelligent), but also we (2)

---

[186] This follows from my non-reduction argument of chapter two.

[187] This is closely related to Daniel Dennett, *The Intentional Stance* (Cambridge, MA: MIT Press 1989) the intentional stance argument, but it does not rely on it. Whether Dennett's argument works or not, it is patently circular to require that a subject possess intentional states in order to deem her intelligent, but to attribute intelligent states to her only after one has judged her to be intelligent.

either exclude animals and young children from the cognitive world or let them in by reasoning in circle.

The fact is that if intentional learning is sufficient for ensuring the intelligence of a *belief*, then we have no good reason to deny that the same sort of learning makes a *process* intelligent as well. No matter what definition of cognition we land on at the end of the day, we can be sure that these central types of states and processes, i.e., those developed as a result of deliberate learning, will remain safely intelligent. As we discover the underlying criteria by which to determine admission into the category of the cognitive, we may end up including many more states, behaviors, and processes than we would admit using intuition alone. But, those procedures central to our case, i.e., those that result from intentionally engaging in a trial-and-error learning process that requires attention, discrimination, control-refinement, etc., are unlikely to be excluded.

Still, we should want some way to discriminate the practical intelligence that results from practice from the other products of repetitive action. After all, there are presumably a multiplicity of changes that occur as a result of repeating an action, but clearly not all these changes will qualify as practical intelligence. The way to differentiate the changes in practical intelligence from other purely physiological or chemical changes, I think, is by counterfactual. We should ask: if the intention to learn was not present in the agent practicing the skill, that is, if the agent was not involved in attempting to improve her performance but was simply repeating certain bodily movements, could the learning still occur? If it is reasonable to imagine changes occurring without the intentional effort or engagement of the agent, then those changes should not be thought of as practical intelligence. Practical intelligence should only describe those changes that result from practice, not from shear exposure or repeated movement. This, I think,

is a good way to differentiate changes like the strengthening of one's muscles from the homing of attention and the refinement of control.[188]

To see how such a differentiation works, imagine a passive subject who is in this condition because she is unconscious, uninvolved or uninterested. Theoretically, we should expect to be able to discriminate, at least roughly, the types of differences in learning that we'd see between her and another subject who was actively involved in the learning process. Let us imagine two individuals performing the same finger motions on a keyboard and creating the same noises as a result of them. Let's say that one is actively engaged in learning to play the piano and the other is having her fingers moved, either by some 3rd person physical manipulation, electrical stimulation or specific instruction, but in the absence of an intention to learn anything as a result of these movements.

What could someone actively engaged in learning learn that someone just moving her fingers would not? The changes that would result in each individual, at least intuitively I think, should differ. After all, we would expect the actively engaged learner to learn things that the passive finger mover would not. Following this line of reasoning, we can at least, in principle, distinguish between the strengthening of the finger muscles that each subject would presumably undergo, and the heightened understanding and control that only the active learner would achieve. The changes that result from intentional engagement such as the refinement of control and attention are constitutive of practical intelligence; the changes that can occur in the absence of intentional practice (i.e., only as a result of repetitive movement) constitute the noncognitive aspects of skill acquisition and instantiation. Of course, which results of learning require an

[188] We may encounter empirical evidence that suggests that intentional action and intentional action for the sake of improving produce identical results. If this turns out the case, then we will, of course, have to rethink what the intention to learn or improve adds to practice that a reason for acting does not.

intention to learn and which do not is an empirical matter. In section 5 of this chapter, I provide empirical evidence that supports this theoretical distinction.

3.1 Practical intelligence as attention and control

In this section, I provide two examples of the practical intelligence of skill. Though it is impossible to present a comprehensive analysis of every effective technique of skill development, I suggest that two common practices operative in skill learning are sufficient for highlighting the practical or embodied intelligence that is characteristic of skill: (1) homing attention and (2) developing control.

When it comes to practicing skills, it is quite common for teachers to direct their students' attention to particular objects, areas, qualities or events. So, it seems that practicing requires focusing one's attention on the right things and developing the habit of noticing relevant features of one's action-space. However, the ability to attend to the right things goes beyond just knowing *that* one ought to attend to a particular location, object, or feature. It even goes beyond knowing *why* one ought to attend to that particular location, object or feature. One must learn *how* to attend to those particular locations, objects and features; this is why practice is necessary.

It seems, then, that as a result of practice, one becomes more sensitive to the relevant features of one's environment. And because picking out features is not usually an end in itself, but rather a prerequisite for a response, practice promotes the ability not only to pick out relevant features but also to load them with significance. That is, increased sensitivity requires drawing connections between relevant properties and appropriate responses. This includes understanding what the relevant features of an environment indicate in terms of the appropriate expectations and predictions to make about them. Additionally, anytime one learns to focus one's attention on the right things, one is simultaneously learning how to filter out others.

All of this, of course, as I am arguing, is not a matter of learning rules, but a matter of learning how to pay attention; it is a matter of increasing one's sensitivity to the relevant features of one's action-space. And the difference here is an important one: one does not become aware of the fact that certain features are important to keep track of (though this may be learned, too) but one learns to keep track of those important features. That is, learning by practice teaches one to attend to the significant features in one's environment and not to the fact that those features are significant. Learning through practice increases one's sensitivity to the relevant features of a perceptual array; it does not necessarily increase one's storeroom of factual information.[189] If it were only the latter, then we would still lack a satisfactory response to the question, "why practice?"

To further reinforce the claim that learning through practice teaches persons to focus their attention in the right way, we have only to look at some recent studies on selective or focal attention. As it turns out, robust evidence suggests that experts are better able than novices at attending to those features that are instrumental to successful skill instantiation. As Zenon Pylyshyn writes,

> In most cases the difference between sports novices and experts is confined to the specific domains in which the experts excel—and there it is usually attributable to the ability to anticipate relevant events. Such anticipation is based, for example, on observing initial segments of the motion of a ball or puck or the opponent's gestures (Abernaty, 1991; Proteau, 1992). Except for a finding of generally better attention-orientation abilities (Castello and Umilta, 1992; Greenfield et al. 1992; Nougier, Ripoll and Stain, 1989) visual expertise in sports, like the expertise found in the Chase and Simon studies of chess skill, appears to be based on the nonvisual abilities related to the learned skills of identifying, predicting and therefore attending to the most relevant places.[190]

[189] See chapter 3, section 3.2 for a discussion of experts and their distinct inability to describe the rules governing what they do.

[190] Pylyshyn, *Seeing and Visualizing* (Cambridge: MIT Press, 2003), 85. Pylyshyn suggests that cognition does not impact perceptual processing, but only the early selection of visual properties. On this issue he and I disagree. But, it is perfectly compatible with my position to concede that early selection is instrumental to skill learning, but also to hold that early selection is not sufficient as an explanation of the diachronic cognitive penetration of early perceptual systems.

Imagine the difference between Roger Federer and well, me, examining an approaching tennis ball. We can both know that we are supposed to "pay attention to the ball," and maybe even to "pay attention to the spin of the ball," but what we are actually able to attend to and what these things mean to us in terms of our readiness to respond to them, couldn't be more different.[191] It is exactly this sort of sensitivity that I claim is an integral aspect of the practical intelligence that is homed through practice and is essential to skill.

3.3 Response, repetition and control

Another trait of practice that is instructive in revealing the nature of practical intelligence is the necessity for repetition, performed not only after verbal instruction, but after physical guidance, adjustment and success as well. That is, it is necessary for one to repeat an action in order to refine one's control, even after one has figured out how to guide one's movements appropriately. When a student successfully performs a skill, the process of skill learning does not come to a close, but rather ensues in the form of repetition. This feature of skill acquisition should alert us to the fact that what is developed in practice involves developing the proper control, timing, force, precision, etc. with which to respond.

What one learns through practice, then, is the appropriate ways in which to instantiate a skill, and not just facts about the proper ways in which to do this (though, of course, one may

[191] For instance, as Wallace, para 21 writes, "This thing about the ball cooperatively hanging there, slowing down, as if susceptible to the Swiss's will — there's real metaphysical truth here. And in the following anecdote. After a July 7 semifinal in which Federer destroyed Jonas Bjorkman — not just beat him, *destroyed* him — and just before a requisite post-match news conference in which Bjorkman, who's friendly with Federer, says he was pleased to "have the best seat in the house" to watch the Swiss "play the nearest to perfection you can play tennis," Federer and Bjorkman are chatting and joking around, and Bjorkman asks him just how unnaturally big the ball was looking to him out there, and Federer confirms that it was "like a bowling ball or basketball." He means it just as a bantery, modest way to make Bjorkman feel better, to confirm that he's surprised by how unusually well he played today; but he's also revealing something about what tennis is like for him. Imagine that you're a person with preternaturally good reflexes and coordination and speed, and that you're playing high-level tennis. Your experience, in play, will not be that you possess phenomenal reflexes and speed; rather, it will seem to you that the tennis ball is quite large and slow-moving, and that you always have plenty of time to hit it. That is, you won't experience anything like the (empirically real) quickness and skill that the live audience, watching tennis balls move so fast they hiss and blur, will attribute to you."

learn facts, too). Importantly, one learns the proper manner in which to move, and what it feels like to do this correctly and, eventually, automatically. One refines one's own movements, by figuring out how to control them properly, by discovering the precise timing, force, position etc., that is required in order to successfully instantiate the complex motor patterns that are constitutive of skill.

One does not learn the fact that *m* is the motor pattern that is required for success, but rather, one learns to instantiate *m*, by learning to control one's movements appropriately. If skill learning was reducible to fact learning then it would be impossible to do justice to the fact that practice and repetition are required for skill acquisition. After all, what explanation could be given for the fact that students are not sent off to ponder the rules governing a skill,[192] but rather, required to perform the skill over and over again in order to learn it? That is, controlling a movement, executing it at the proper time, with the appropriate force, with the appropriate weight distribution, body angle etc., is not a matter of simply knowing the names for specific physical movements. While a coach may be able to yell "now" or "harder" and, thus, know when or with how much force to execute a skill, she may still lack the ability to perform the skill successfully. And in order to learn *that*, practice is required.

Importantly, though what is learned is not a fact, it can still be right or wrong. That is, one can misjudge the proper force required to throw a baseball over home plate, throwing it too hard or not hard enough. One can misjudge where one's leg is located in space when one is attempting to kick a soccer ball and one can slip as a result of turning too quickly on one's foot. Though these are in fact mistakes, they are not mistakes about facts; they are mistakes in the judgment and control required for integrating inputs and outputs into successful, intentional

[192] It is interesting to note, however, that visualization is often a part of skill learning. I don't think this means that skills are concepts, but it certainly does indicate a close relationship between thinking and acting skillfully.

action. That is, actions can be normative as can beliefs, but the former aim at success whereas the latter at truth. If these were factual mistakes, then correction could be achieved as easily through information collection as it is by practice. But this is clearly not the case.

3.4 Objection

At this point, one may claim that it is not the case that the nature of the things learned through practice are different from facts that are learned through, say, communication, but rather, that we simply do not have the proper vocabulary to make the types of fine-grained distinctions that are necessary for communicating the kinds of things that are learned as a result of practice. That is, one may claim that we need to practice because we don't have the words to express the facts that one learns in practice. After all, it's easier to show someone a picture than describe it, and it's easier to present an example than to give a definition. One could claim that when we learn how to e.g., play a scale on a piano, all we learn is the fact that "*x*, *y* &*z*" is the appropriate way to play that scale"[193] This objection holds that there is nothing practical or embodied about this kind of intelligence; it is just another example of good old-fashioned propositional knowledge.

Unfortunately, this line of defense fails to account for the fact that repetition of a skill is required long after one becomes familiar with the fine-grained but "ineffable" qualities that one may come to know through experience. That is, if what is learned through practice is just the fact that to perform *a*, *S* must do *f*, where *f* is a fact that is inexpressible in language, then it is far from clear what we would gain by repeating the process by which *S* came to know *f*. If that which is learned through practice is a fact, then reflecting on it should be as effective as repeating it. After all, when learning facts, we are often encouraged to go over them in our mind.

[193] This is essentially the Jason Stanley and Timothy Williamson, "Knowing How," *Journal of Philosophy* 98 (2001): 430 position. For objections to this view, see chapter three, section 2.

If what was learned through skill was of this nature, then one should be sent home to contemplate the qualities of *f*. But this is obviously not how we improve in skill. We improve our skills by repetition, and not just through reflection (though, we surely may need to reflect as well). The repetition required for skill learning indicates that what is learned through practice is not just factual knowledge, but the practical judgment and control necessary to successfully perform a skill. Repetition through practice develops one's responsiveness, not one's awareness of facts about appropriate responses. This type of learning, that is, the development of the appropriate responses and controlled movements required to instantiate complex motor patterns, is what I call practical intelligence.

The conflation between learning how to respond and learning a fact about the right way to respond parallels, I think, the conflation between conscious experience and introspective experience. It is tempting to say that any experience of *q* is an experience of oneself as experiencing *q*. However, this constitutes a conflation between introspective consciousness and transitive consciousness.[194] This position not only leads to an infinite regress, but also gets the phenomenology wrong. Often I am concerned only with the world, and not with myself as experiencing the world. Likewise, when it comes to skill, I learn how to perform a skill, and not a fact about my performing that skill. The learning is at the conscious motor level, but not necessarily at the introspective one. The temptation to explain skill acquisition in higher-order terms, I think, stems from the fact that we have a fairly sophisticated way of dealing with the logic of propositions, but not with the logic of actions. But, just because we do not possess a

[194] See David M. Rosenthal, "Two Concepts of Consciousness," in *The Nature of Mind*, ed. David M. Rosenthal, 466 (New York: Oxford University Press, 1991); David M. Rosenthal, "State Consciousness and Transitive Consciousness," *Consciousness and Cognition* 2 (1994): 361-2; David M. Rosenthal, "Being Conscious of Ourselves," *The Monist* 87 (2004):164.

logical framework to analyze know-how, does not mean that a propositional framework can provide us with an adequate account of skill.

In this most basic sense then, practice is the key to developing an increased sensitivity to the relevant features of an environmental array. Practice develops a responsiveness that enables one to make fine-grained distinctions and appropriate, precise, effective, and expedient responses to those discriminations. Practice leads to practical intelligence in the form of a mastery of the inputs and control of the outputs constitutive of the motor patterns that express successful skill instantiation. It teaches one to pay attention to the right things and respond appropriately to what one notices. This type of learning is clearly intelligent since it involves attention, discrimination, attributing significance to relevant information, trial-and-error, and the capacity to revise responses according to the success or failure of particular trials.

3.5 Is practice necessary for developing practical intelligence?

Of course, we should want to know whether practical intelligence must be developed through practice, or if it is possible to possess practical intelligence without ever having had to learn it. We can ask, like Donald Davidson does about mental content,[195] if a man was to emerge from a swamp fully formed, with a great many abilities, would those abilities be skills? That is, could swamp-skill be characterized by practical intelligence?

The short answer, is "no." Though one may perform actions that are indistinguishable from skills, not having developed those abilities through a process of practice precludes them from being characterized by practical intelligence.

Largely, this is definitional, but I think there are reasonable considerations supporting this position. The main one is that practical intelligence characterizes skills, and not the

[195] See Donald Davidson, "Knowing One's Own Mind," *Proceedings and Addresses of the American Philosophical Association* 60 (1987): 441-458.

instantiations of skills. So, though it is possible to have identical extensional instantiations of an action, what those actions are instantiations of, may differ. In this case, they would differ, of course, in their history. But it is exactly their history that is responsible for developing the cognitive character of a skill. This is because, as Fred Dretske has argued, learning harnesses inputs to outputs and makes behavior meaningful.[196] Learning is the initiator of intelligence.

Now, an unlearned ability may be sophisticated and impressive, but there is no reason to believe that it is not brute or instinctive.[197] There is no way to gauge that the ability expresses features that are flexible, dynamic, deliberate and meaningful. The only way to determine this is if the ability is responsive and sensitive to environmental features and goals; that is, if it learns.

My main claim here is that because we can distinguish between the practical intelligence of a skill, and the expression of that intelligence in a skill instantiation, there is no reason to believe that swamp-skill matches real skill in anything but its expression.

4. Automaticity and intelligence

In my second chapter's discussion of automaticity and cognition, I argue that being automatic does not entail being noncognitive or unintelligent. This is because automaticity on the level of information processing does not entail automaticity on the level of rule implementation. So, we can imagine a particular instantiation of a skill being performed in quite automatic fashion, without thereby making the procedures that are followed by that instantiation insensitive or unresponsive to various environmental features or goals, experience and learning.

After all, let us consider the following case: if through rigorous practice we became able to do arithmetic automatically (by adding numbers without going through any conscious or deliberative strategy), it would be rather strange to call this process noncognitive. This seems

[196] Dretske, *Explaining Behavior*, 104.
[197] Think, for instance, of the dance of a honeybee.

true, especially because the way in which arithmetic becomes automatic requires a regular dose of deliberate, cognitive sophistication, practice and complexity. As such, the automatic quality of some particular arithmetic calculation would hardly undermine the cognitive nature of the arithmetic procedures that are instantiated therein.

The same should be the case for hitting a tennis ball, doing a handstand, or playing a combination of notes. Since the procedures that are followed in order to perform these skills are structured and implemented by intelligent processes, the particular automatic performance that is instantiated according to those procedures need not undermine the overall intelligence of a skill.

Importantly, we should note that since a skill is a complex of motor patterns that is developed over time and instantiated at different times, we should avoid identifying skill with any one of its particular instantiations. The locus of practical intelligence is not an instantiation of a skill, but the procedures that constitute the skill. So, a skill instantiation is characterized by practical intelligence not because it is conscious or deliberate or flexible, but rather because it is the performance of a particular set of procedures that have been developed through practice. That is, skill instantiations inherit their intelligence from the procedures that they are instantiations of.

5. The relationship between skill and perception

At this point, I'd like to revisit the definition of cognitive penetration that I developed in chapter two. In this way, I can elucidate what must be true of the relationship between skill and perception if it is to qualify as instance of cognitive penetration.

Earlier, I define cognitive penetration as changes in perceptual processing that occur as the result of a logical relationship to the meaning or content of a cognitive state.[198] So, the question to ask in order to determine the presence of cognitive penetration is, "Is the meaning or

[198] See chapter two, section 3.2.

content of a cognitive state explanatorily relevant to the functioning of an early perceptual system?" If the answer is "yes," then the perceptual system is cognitively penetrable.

Recall that the reason explanatory relevance is required for cognitive penetration, rather than any old relation, is because we should not want any impact from any feature of a cognitive state to count as cognitive penetration. We should want to avoid the rather odd conclusion that when I exercise and my metabolism speeds up, this entails that my metabolism is cognitively penetrated by the intentional states that guide my exercising behavior. We should want the cognitive element of a state to impact the processing of a system in order for it to constitute an instance of cognitive penetration.

Now, it is natural to talk of the meaning or content of an intentional state, but since embodied skills are not intentional states but procedures, we must find a satisfactory way to talk about the cognitive aspect of them. In order to transfer the notion of content and meaning from states to processes, it helps to clarify the role that semantic content plays in the constitution of a representational state. Quite clearly, it is the semantic or meaningful content of a representational state that makes that state cognitive. That is, there are various aspects of representational states—some physical, some structural, and some meaningful—it is only the meaningful aspects of a representational state that make intentional content even potentially intelligent. If the content of a representational state was not meaningful, if it had no semantic value, then that state would not qualify as cognitive.[199] As such, if we can find that aspect of skill which makes the instantiation of complex motor patterns not just processes, but processes

[199] Following Fred Drestke, *Naturalizing the Mind* (Cambridge, MA: MIT Press, 1997), 4 a state may be representational, in that it may carry information, but if it is not meaningful information, that is, if it has no semantic value for the agent or system, then that state is not a cognitive state. There is a difference between carrying information and having semantic content.

that are characterized by cognition, then this will be the aspect of skill that, if related in a logical way to perceptual processing, makes for an instance of cognitive penetration.

As I've argued in the preceding paragraphs, what makes a skill cognitive is the practical intelligence which characterizes it. That is, it is the particular style of sensitivity and control that is developed through practice which constitutes the cognitive aspect of a skill. Therefore, if it is this aspect of skill, its practical intelligence, that expresses a logical relationship to perceptual processing, then this constitutes an instance of cognitive penetration. Quite simply, if the function of a perceptual system cannot be explained without appeal to the practical intelligence of skill then a logical connection between cognition and perception has obtained. But if the particular manner in which a skill is instantiated is irrelevant to explaining early perceptual processing, then an instance of cognitive penetration is absent.

Let us take an example of an adult learning to play the piano in order to help clarify this relationship. Part of what goes into learning to play a piano is figuring out which keys to strike, how hard, and when. These types of refinements are based on listening and responding to the sound of the notes that one produces when one strikes various keys in various ways. One cannot learn to play the piano without actively engaging in a learning process. So, when instantiating the skill of playing the piano, one plays in a particular fashion; one attends and responds to with a particular sensitivity and control that one has developed through practice. Now, if instantiating this skill produces changes in early auditory processing such as e.g., an increase in the processing that is allocated to chords,[200] and if this change cannot be explained without talk of the manner in which the piano is played, that is, if it cannot be explained by reference to finger movement or being in the presence of music, but rather requires appeal to the particular attention and response

[200] We can assume that such processing is correlated with a heightened capacity to discriminate the individual notes that comprise chords.

that is homed through practice, that is, if the explanation requires reference to practical intelligence, then this type of relationship is one of cognitive penetration.

Additionally, such cognitive penetration would be diachronic, since both the development of practical intelligence and the regular instantiation of a skill, cannot happen all at once. The types of changes in early perceptual processing that may result from skill, would be the consequence of playing the piano over a period of time, and thus, instantiating practical intelligence at different times. This makes it clear that changes in perceptual processing are, by nature, conservative; they are not the result of whim or fancy, but rather, of the regular and successful instantiation of practical intelligence in action.

So, if (1) there are processing changes in an early perception system, and (2) if in order to explain those changes one must appeal to the practical intelligence of a skill, then, (3) this is a legitimate case of cognitive penetration.

In order to constitute a legitimate instance of cognitive penetration, there must be changes in early perceptual processing that cannot be explained without appeal to practical intelligence. That is, only an appeal to the sensitivity and control developed through intentional practice should be able to give an account of why particular changes in perceptual processing occur. As such, explaining the impacts of skill on the function of perceptual processing will require talk of learning, attention, sensitivity, response, control, discipline, etc. It will require talk of practical intelligence. However, as I mention above, any cognitive state or process will have noncognitive counterparts. Therefore, in order to determine if some relationship is a legitimate instance of cognitive penetration, we must have a way to distinguish the cognitive from the noncognitive features of a skill. That is, we must be able to distinguish the practical intelligence that characterizes a movement from the movement itself.

In section 3 of this chapter, I propose a theoretical solution for distinguishing between practical intelligence and the noncogntive aspects of a skill: we do so by counterfactually determining if that which is learned as a result of practice must have been developed in the presence of an intention to learn it. Here, I would like to gesture to some empirical findings that seem to warrant this kind of distinction. This empirical evidence supports the idea that the practical intelligence of skill can be distinguished from its other features. Naturally, this reinforces the intuition that what is learned when one is actively involved in learning a skill is different from the changes that occur if one is simply performing motions. The following empirical evidence gives us *prima facie* reason to believe that a distinction between the results of actively engaging in learning, as opposed to simply acting, can be drawn. Researchers have found that rats which were involved in learning developed a denser network of synapses per nerve cell than their motorically active but not cognitively engaged counterparts (i.e., the rats that were simply moving around). [201] The experiment progressed as follows:

> There were four groups [of rats] in all. One group of rats was taught to traverse an elevated obstacle course; these "acrobats" became very good at the task over a month or so of practice. A second group of "mandatory exercisers" was put on a treadmill once a day, where they ran for 30 minutes. A third group of "voluntary exercisers" had free access to an activity wheel attached directly to their case, which they used often. A control group of "cage potato" rats had no exercise. What happened to the volume of blood vessels and number of synapses per neuron in the rats? Both the mandatory exercisers and voluntary exercisers showed higher densities of blood vessels than either the cage potato rats or the acrobats, who learned skills that did not involve significant amounts of activity. But when the number of synapses per nerve cell was measured, the acrobats were the standout group. Learning adds synapses; exercise does not. Synapse formation and blood vessel formation (vascularization) are two important forms of brain adaptation, but they are driven by different physiological and by different behavioral events.[202]

---

[201] J.E. Black, K.R. Isaacs, B.J. Anderson, A.A. Alcantara and W.T. Greenough, "Learning Causes Synaptogenesis, Whereas Motor Activity Cause Angiogenesis in Cerebellar Cortex of Adult Rats," *Proceedings of National Academy of Sciences MSA* 87 (1990): 55568-5572.

[202] John D. Bransford, Ann L. Brown, and Rodney R. Cocking, *How People Learn: Brain, Mind, Experience and School* (Washington D.C.: National Academy Press, 1999), 107.

Such findings,[203] give us good preliminary reason to hold that the neurological changes that result from learning can be distinguished from those that result from movement and action where a deliberate intention to learn is absent. That is, such findings reinforce the plausibility that we can differentiate those changes in perceptual processing that are caused by physical movement, from those changes that are cognitively driven. So, both conceptually and neurologically, it is possible to differentiate practical intelligence that results from learning a new skill intentionally from the affects of movements that do not involve cognitive elements such as attention, awareness, control, problem-solving, etc. In this way, potentially, we have the opportunity to determine, not just theoretically, but also empirically, when the cognitive aspect of a skill is logically related to the function of the early perceptual modules.

Of course, this is not an argument for the fact that cognitive penetration occurs; rather, it is an explanation of the relationship between skill and perception that would need to obtain in order for said relationship to constitute a legitimate instance of cognitive penetration. To be clear, it is not my ultimate goal to prove that cognitive penetration occurs as a result of the interaction of skill and perception. Rather, in pure philosophical style, I present the conditions that must be met in order for some relationship to qualify as an instance of the cognitive penetration of perception by skill. I leave filling in the empirical evidence to the scientists.

---

[203] Mark R. Rosenzweig and Edward L. Bennett, "Cerebral Changes in Rats Exposed Individually to an Enriched Environment," *Journal of Comparative and Physiological Psychology* 80 (1972): 304-313; F. Gonzalez-Lima, P. A. Ferchmin, V. A. Eterovic and E. M. Gonzalez-Lima, "Metabolic Activation of the Brain of Young Rats after Exposure to Environmental Complexity," *Developmental Psychobiology* 27 (1993): 343-51.

CHAPTER 6

1. Conclusion
1.1 Applications: Enactive Perception

In the introductory chapter of this dissertation, I stated that my starting off point for thinking about skills and their relationship to perception began with the theory of enactive perception. Now, in light of my investigation into the nature of skill and cognitive penetrabilty, I'd like to reconsider the role that sensorimotor skill plays in perception. Quite clearly, the theory that I have detailed for skills, generally, should apply to sensorimotor skills, in particular.

I do not intend to haggle over whether early perception *is* impacted as a result of diachronic influences. Rather, as I have throughout this dissertation, I assume that changes in early perceptual processing may in fact occur.[204] In this section, specifically, I assume that there is a possibility that changes in early perceptual processing are caused by sensorimotor knowledge.[205] What I'd like to do here, then, is to examie if the relationship between sensorimotor skill and perception constitutes an instance of cognitive penetration.

In this conclusion, I intend to assess if sensorimotor skills can meet the three criteria for skillhood. By applying this test, I can determine if sensorimotor skills are characterized by intentionality, success and practical intelligence and, as such, whether they can stand on the cognition side of the perception-cognition relationship, which constitutes cognitive penetration.

---

[204] As I have reiterated, this move is justified because even the opponents of cognitive penetration admit that changes in early perceptual processing occur, though they deny that these changes are instances of cognitive penetration.

[205] See introduction for empirical evidence justifying this assumption. See also Richard Held and Alan Hein, "Movement-Produced Stimulation in the Development of Visually Guided Behavior," *Journal of Comparative and Physiological Psychology* 56 (1963): 872-876.

To begin, we should ask: when instantiated, do sensorimotor skills constitute intentional actions? That is, can sensorimotor skills meet criterion (1)? The answer to this question, I think, will depend on the kind of sensorimotor skills we consider. Sometimes it is easy to identify the pro attitude and instrumental belief, which causes the instantiation of a sensorimotor skill. At other times, however, the instantiation of a sensorimotor skill does not seem to be caused by any reason whatsoever. This does not make the instantiation random or involuntary, [206] but it does preclude it from being intentional.

Some instantiations of sensorimotor skills are overtly related to reasons. For instance, you may ask me to help you find the earring that you lost on the floor. In order to do this, you may instruct me to pass my eyes over each parquet tile in a counterclockwise fashion, moving one step to the right after completing each square (perhaps you read an article about search techniques and this was heralded as the most effective one). In this case, I move my eyes and my foot for reasons, namely because I want to help you find your earring and because I believe that you are right about the best way in which to accomplish this task. Importantly, I not only search the floor intentionally, but I also move my eyes in a counterclockwise fashion and my foot one step to the right, intentionally as well.

There are other sensorimotor skills, however, that are voluntarily instantiated, but not instantiated for reasons. Think of eye movements and gaze fixations here. I frequently scan a room, fixate on particular objects and saccade my eyes without doing so for reasons. The fact is, frequently, I just look around. Contrast this to reading, or admiring a painting, or searching for a lost earring, or trying to find an acquaintance at a party. The upshot here is that if looking often fails to be an intentional action, then it would be very strange indeed to regularly find reasons for

[206] See J.L. Austin, "A Plea for Excuses," in *Classics of Analytic Philosophy*, ed. Robert Ammerman, 379-398 (USA: Hacket Publishing, 1990).

eye saccading or gaze-fixing. On the enactive view, however, these activities qualify as sensorimotor skills. [207]

Alva Noë states that "according to the enactive approach, vision depends on one's knowledge of the sensory effects of, say, eye movements, for example, movements of the eye to the right causes a shift to the left in the retinal image." [208] Of course, this does not preclude that eye gaze may, in fact, be intentional. You may, for example, ask me to turn my gaze in order to point out the way that the room shifts as I do this. In this case, I gaze intentionally. But this is very different from what I usually do when I just look around. That is, I may look for reasons, and I may even fixate my gaze for reasons, but more often than not, I do not perform these actions intentionally. However, if these very activities are candidates for skills, then they fail to meet criterion (1).

The enactive theorist may counter with two responses. First, she can accept criterion (1) and, consequently, edit her list of sensorimotor skills to meet criterion (1)'s requirement of intentionality.[209] Alternately, the enactive theorist can reject the proposal that criterion (1) is necessary for skill. She could argue that being done for reasons is too high a bar to set for intentional action. Perhaps, just being voluntary or goal-intended could suffice—in this way, tacit reasons, or implicit reasons or pseudo-reasons could satisfy an amended version of criterion (1). This sort of account would, of course, need to be capable of sustaining a robust notion of

[207] An obvious question that arises here is at what point does an action cease to express a sensorimotor skill. That is, walking is quite obviously a sensorimotor skill, but what about swinging one's arms, or adjusting one's toes to the earth beneath one's foot? Are these each individually sensorimotor skills? Visual search can be intentional, and gaze fixing can be intentional as well, but what about pupil dilation? It is quite clearly necessary for proper visual perception, and it is responsive to light levels, but is this a sensorimotor skill, too?

[208] Alva Noë, *Action in Perception* (Cambridge, MA: MIT Press, 2004), 25.

[209] One may reasonably change one's vocabulary in order to distinguish between these two different kinds of sensorimotor activity. For instance, one may want to reserve "skill" for those sensorimotor capacities that meet criterion (1) and e.g. call those that do not meet the intentional action criterion, sensorimotor abilities.

intentionality and agency, and defending this claim is where the difficult work for an enactive theorist will lie.

Moving on, it appears that sensorimotor skills should have no trouble satisfying the ability criterion or criterion (2). After all, sensorimotor skill is the expression of the mastery of sensorimotor contingencies. That is, successfully instantiating practical actions is inherent to the embodiment of sensorimotor knowledge. In fact, there would be no talk of sensorimotor skills, if creatures were not able to successfully navigate their environments. As such, ability is the basis of ascriptions of sensorimotor knowledge. Therefore, it would seem that in the absence of ability, there would be no temptation to ascribe possession of sensorimotor skills. Sensorimotor skills, then, can easily satisfy criterion (2).

Additionally, like other skills, the ability to instantiate sensorimotor skills is not required under all circumstances, but only under normal counterfactual conditions. So, we should not expect a person to experience the visual presence of the backside of a vase in the dark, or to predict the location of a sound in a quickly rotating room. But the inability to successfully instantiate sensorimotor skill in unusual environments does not count against the possession of sensorimotor skill. Just so long as an action can be instantiated successfully under normal circumstances, criterion (2) is satisfied.

Next, let us consider whether sensorimotor skills satisfy criterion (3), the learning criterion for skill. We should ask, are sensory motor skills learned through practice, where practice is defined as repeated action performed for the sake of improvement? The answer to this question mimics the problem that sensorimotor skills faced with criterion (1). That is, the answer to this question depends on which sensorimotor skills we consider.

For instance, if we take the inverted glasses experiment, it is quite clear that subjects are intentionally trying to learn sensorimotor skills.[210] In fact, part of the experimental design requires subjects to navigate their environment in hopes that they will adapt to their new visual environment. In this way, subjects are engaged in trying to develop motor skills and coordination in order to overcome the perceptual confusion caused by the inverted lenses. It is clear then, that these subjects develop sensorimotor skills as the result of practice.

Another clear example of a sensorimotor skill being developed through practice is illustrated by the disembodied lady, Christina, whose efforts to overcome her proprioceptive deficiencies are documented by Oliver Sacks.[211] We can see that Christina has the intention to improve her sensorimotor abilities since, in the absence of proprioceptive information, she explicitly focuses on guiding her movements using visual feedback. Chistina actively attempts to learn to use visual stimuli in order to compensate for the lack of her proprioceptive inputs. This same phenomenon is observable in cases of stroke victims who intentionally set out to re-learn basic motor tasks like walking, reaching and grasping by repeating these tasks until they can successfully instantiate them. Quite clearly then, the above scenarios can satisfy criterion (3).

Before moving on to consider the types of abilities that are not learned as the result of practice, it is important to recall that according to my definition of practice, the intention to learn or to improve need not correspond to that procedure that is learned as a result of the practice. So, one may practice in order to learn to jump rope, but instead may learn to hop on one foot in a

---

[210] See chapter one, section 1 for a description of this experiment.

[211] Oliver Sacks, "The Disembodied Lady," in *The Man Who Mistook His Wife For A Hat And Other Clinical Tales* (New York: Touchstone Books, 1985), 43-54. See also Jonathan Cole and Jacques Paillard, "Living Without Touch and Peripheral Information about Body Position and Movement: Studies with Deafferented Subjects," in *The Body and the Self*, ed. Jose Bermudez, Anthony Marcel and Naomi Eilan, 245-66 (Cambridge, MA: The MIT Press, 1995).

straight line. The hopping, on my account, would qualify as an ability that has been acquired through practice.

Examples of this kind of mismatched learning can be found in experiments on implicit learning.[212] The most famous of these experiments employs an artificial grammar in order to show that subjects learn grammars without knowing what they are learning and without being able to state the rules governing their newly acquired knowledge. What is important for my purposes, however, is that while subjects are not trying to learn an artificial grammar, they are engaged in learning a seemingly unrelated task. "[S]ubjects were not informed that they were working with rule-governed stimuli. They were merely requested to memorize strings of letters in what was touted as a rote memory experiment." [213] Clearly, memorizing strings of letters involves the intention to memorize these strings of letters. So, subjects have reasons for memorizing letters, but instead they learn a synthetic grammar—what subjects actually learn does not correspond to the reasons for which they practice. On my definition of practice, however, implicit learning satisfies criterion (3). The importance of this here is for us to note that if sensorimotor skills have this mismatched character, they can still qualify as being learned as the result of practice.

Unfortunately, however, it does not seem that practice can account for all sensorimotor learning whether with a corresponding or mismatched intention. Rather, it seems likely that

---

[212] It is not often noted, in discussions of implicit learning, that subjects are actively engaged with a set task. They are not passive subjects, but rather, they follow a given set of rules and instructions. They are directed to learn *something.* This, I think, is a very important point to consider when thinking about implicit learning and its implications.

[213] Arthur S. Reber, "Implicit Learning and Tacit Knowledge*," Journal of Experimental Psychology: General* 118(1989), 221; D.C. Berry and D.E. Broadbent, "On the Relationship Between Task Performance and Associated Verbalizable Knowledge," *Quarterly Journal of Experimental Psychology* 36 (1984): 209-231; D.E. Broadbent and B. Aston, "Human Control of a Simulated Economic System," *Ergonomics* 21 (1978), 1035-1043; Arthur S. Reber and R.B. Millward . "Event Tracking in Probability Learning," *American Journal of Psychology* 84 (1971): 85-99; P. Lewicki, T. Hill and E. Bizot, "Acquisition of Procedural Knowledge about a Pattern of Stimuli that Cannot be Articulated," *Cognitive Psychology* 20 (1988): 24-37.

some sensorimotor skills are developed in the absence of any intention to learn or to improve whatsoever. For example, it would be patently odd to ascribe to an infant who is reaching for a hanging toy, the intention to reach in order to learn or to improve. While the infant may reach and grasp *intentionally* (depending on the age of the infant), it just seems implausible that the reason for her reaching and grasping is to learn or to improve at anything, never mind to improve her motor coordination. It should be clear, however, that reaching and grasping behavior produces motor coordination and sensorimotor skill. The fact is that in many cases where children learn through play, it is unclear whether any intention to learn is causing their behavior; it seems just as likely that they are just engaged in play and happen to learn things, including sensorimotor skills, as a result. As such, sensorimotor skill learning of this sort will not be guided by an intention to learn and thus will not satisfy criterion (3).[214]

Of course, this is all an empirical matter. It is quite possible that an intention to learn is always present in the development of senorimotor skills. In the absence of appropriate evidence, however, it is still useful to map out the options that an enactive theorist has *if* it does turn out that sensorimotor skills are not learned through practice.

The first option is to accept criterion (3) as it stands and assert that only certain sensorimotor skills qualify as genuine skills. The second option is to refashion my definition of practice. The enactive theorist may argue that an intention to learn is not required for a skill to be learned through practice, but instead, may insist that an intention *of any kind* is sufficient to sustain the notion of practice. In order to ensure that most sensorimotor skills will satisfy the learning criterion, the enactive theorist could argue that practice should be defined as repeated intentional activity and not repeated intentional activity for the sake of learning. A good way to

[214] See Alison Gopnik, "Your Baby is Smarter Than You Think," New York Times, August 15, 2009, Opinion, New York edition.

support this claim would be to provide empirical evidence suggesting that the intention to learn does not contribute much to the learning process over and above what intentional performance provides. If such evidence were forthcoming, then one could claim, in a principled way, that the difference between intentional action, and intentional action for the sake of learning, is negligible. As far as I can tell, this remains a relatively open question.[215]

2. Implications:

2.1. Epistemology

Traditionally, foundationalist epistemologies of the empiricist variety have been committed to the idea that directly experienced perceptual states can act as a secure foundation for knowledge.[216] According to them, in order for knowledge to have a solid foundation, there must be some stage of belief formation that is necessarily veridical. Sensory experience, it is supposed, can play this role since sensory states are known immediately—what makes them true is right there in the experience of them. Thus, directly experienced states contain no room for error; they are always veridical and as such, a secure foundation for knowledge.

Unfortunately, the necessary veridicality of directly experienced perceptual states would be compromised if these experiences were impacted by beliefs, desires and thoughts. After all, beliefs and thoughts can be true or false, and if direct sensory experience were tainted by them, then direct sensory experience could be true or false, as well. But this possibility of error undermines the foundation of knowledge that direct acquaintance is supposed to provide. For this reason, foundationalist epistemologists of the empiricist bent adamantly reject the possibility of interference by cognition with perception.

---

[215]See chapter five, section 5. Did the rats have an intention to learn? At the very least the learning did something that action alone did not, but was it correlated to an intention to learn or just intentional action? This is the important question that an enactive theorist would have to confront here.

[216] "Perhaps the best known proponent of an acquaintance theory is Bertrand Russell, but it takes little imagination to read the view into most of the British empiricists" Stanford Encyclopedia of Philosophy, "Foundationalist Theories of Epistemic Justification," http://plato.stanford.edu/entries/justep-foundational/, section 2.2, para 4.

In contrast to the threat that cognitive penetration by propositional states has on the possibility of knowledge, cognitive penetration by skill has a disarming affect on those who have argued that the impacts of cognition on perception would guarantee the reign of skepticism. Unlike beliefs, which may be ubiquitously false, except for the one about my own existence, skills have a reality-safeguard built into them via the success criterion. After all, if one was sensitive to the wrong things, and constantly misinterpreted the relevance of environmental stimuli, and responded inappropriately to the wrong features of a perceptual array, then one simply would not develop skills. Skill acquisition requires regular success and feedback. This means that a reliable connection between skill and reality is built into the possession of a skill.

Skills have a safeguard against radical skepticism because skills cannot succeed with a radically distorted relationship to the world. As such, if skills were to impact perception, they could not do so in a haphazard way. A being with false beliefs may be a fool, but a being without skills would be dead. In this way, the threat of erroneous cognitive penetration by skill is naturally circumvented.

Happily, it not only turns out that the diachronic cognitive penetration of perception by skill does not threaten the possibility of knowledge, but in fact, it cultivates it. This is because the better one gets at interacting with the world, the more appropriate and fine-grained one's sensitivity to relevant stimuli becomes and the better one is at manipulating and responding to those stimuli. On my theory then, the cognitive penetration of perception by skill brings one closer to the world, not farther away from it. The better we get at skill the more likely it is that we are encountering the fine-grained qualities of an action-space. If skill impacts perception, then, it would likely make perception more sensitive to the richness and fullness of one's environment.

This result, I think, is highly desirable. As opposed to traditional epistemologies that must eschew learning in order to locate the foundations of knowledge, my account provides a way to approach learning as an instrument for discovering truth. That is, what we learn does not take us further away from certainty, but rather, it brings us closer to it. And this just seems right.

2.3 Evolutionary biology

Zenon Pylyshyn, in citing the reasons against the cognitive penetrability of perception, claims that perceptual processing, if it were influenced by thoughts and expectations, would ultimately make a creature less viable.[217] What he means is that seeing what you want to see, rather than what's really there in front of you, is dangerous. After all, an animal never expects to perceive the unexpected and if perceiving is dictated by expectations then perception could never inform a creature of anomalous dangers. This would quite clearly present an obstacle to a creature's survival since survival often hinges on a creature's ability to respond to unexpected situations. Cognitive penetrability, as Pylyshyn argues, entails that a creature responds to her own impressions, desires and expectations of reality rather than to reality itself.

Cognitive penetration, construed in this way, would liken perception to imagination and thus, undermine its most useful evolutionary function: to inform. Distancing perception from detection or monitoring would quite clearly have little practical advantage for an organism.

As opposed to the free-for-all of cognitive penetration by occurent propositional states, however, the gradual and controlled type of cognitive penetration, which may result from the diachronic influences of skill, would not be susceptible to the above criticism. This is because the types of changes that happen over time and are tied to success in the way that skills are, are necessarily conservative in nature. The risk that perception, if cognitively penetrable, would

[217] See Jerry Fodor, *The Modularity of Mind: An Essay on Faculty Psychology* (Cambridge: MIT Press, 1983), 68.

become a sort of haphazard imagining instead of an accurate informing does not arise here since skills are necessarily tied to experience and success in a way that grounds them in reality.

Moreover, there are evolutionary advantages to *not* having an absolutely rigid perceptual processing system. There is an advantage to not having to consult long-term memory and other global processes in order to cope with perceptual regularities. By internalizing constant features of environmental arrays, a creature can react to familiar stimuli in an efficient manner while freeing up the mental space to encounter novel and complex events. Such a creature would be better able to react to unexpected dangers since it would not be preoccupied with higher-order processing of regularly encountered perceptual inputs. The fact is that there is nothing contradictory in a perceptual system that is influenced by learning and experience but also sensitive to novel stimuli. This is presumably how the learning happens. And such learning has distinct evolutionary advantages over inflexible processing.

After all, in terms of the evolution of perceptual systems there are only two options: either perceptual systems evolve by accident, or perceptual systems are equipped with mechanisms by which to adapt to a creature's environment in relevant and advantageous ways. Of the above options, the kind of changes in perceptual processing that are linked to the activities and needs of the organism and not simply the results of chance mutations, are clearly the more superior. As such, diachronic cognitive penetration of the conservative variety that I have described, would seem to be the one most favored by evolution. That is, diachronic cognitive penetrability of perception by skill can preserve the benefits of a conservative perceptual system, while gaining the advantages of intra-generational adaptation. Of course, this doesn't mean that this is the way perceptual processing works, but it certainly is as good an argument for cognitive penetration as Pylyshyn's is against it.

2.3 Naturalism and Dualism:

The idea that mind and body compose two ontologically distinct realms has plagued philosophy since its inception.[218] The dualism of traditional theories of perception, which hold that intentional and qualitative character constitute independent elements of a perceptual experience, is hardly subtle in its continuation of this tradition. It is often forwarded that the qualitative character of a perceptual event is the physical or bodily aspect of perceptual experience; it is mechanistic, hard-wired and absolutely immune to the impacts of cognition. The intentional aspect of perception, on the other hand, belongs to the mental realm. This is where thoughts, judgments and categorizations of sensations arise, develop and change. The separation of these two components of perception retains the spirit of dualism and the difficulty associated with their interaction propagates mystery.

The possibility that perception is diachronically impacted by skill undermines the dualist tendency of philosophy in two ways: first, the idea that intelligence is a property of actions, not reducible to thoughts about or intentions for those action, solves the problem of dualism by not letting it begin. That is, because intelligence is not exclusively a property of representational states, the mystery of dualism cannot get off the ground. Secondly, the idea that embodied intelligence can impact the perceptual processing of sensory inputs offers a straightforward interaction between physical processing and intelligence. Since the practical intelligence of skill is embodied intelligence it does not inspire confusion as to how one bodily feature can interact with another. The interaction of two bodily processes is hardly a conundrum.

Further, the practical intelligence of skills emphasizes a convergence between the mental and physical rather than focusing on the incompatible characteristics of brute mechanical

[218] Though Descartes is blamed for propagating dualism, Plato had already settled on the fundamental separation between mind and body as a basic tenet of Western philosophy many years before.

function and reasoning. By looking at embodied intelligence, we begin at the overlap between mental and physical rather than appealing to phenomena that seem disembodied because we do not yet have the brain science to unify them. By beginning with practical actions that are aimed at success instead of with propositional knowledge that is aimed at truth, we have, at the very least, the possibility of developing a unified account of cognition. This is no guarantee, of course, but it does seem that the opposite strategy has failed, so the promise of a new direction is welcome.

Additionally, a satisfactory account of the practical intelligence of embodied actions opens up new and natural ways in which to approach the intelligence of several domains of thought for which a propositional analysis of cognition does not suffice. For example, a skill-based account of intelligence offers an alternative for thinking about ethical and aesthetic judgments. After all, much of our ethical reasoning simply cannot be described as the application of the categorical imperative or as the result of a utility calculus.[219] But simply because a straightforward analysis of ethical reasoning cannot be reduced to truth-functional propositional reasoning does not mean that ethical judgments must be brute or noncognitive.[220] Importantly, by applying a skill-based account to ethics, we can explain how a habit-centric virtue ethics can remain firmly cognitive. That is, my account of the practical intelligence of skill opens up the possibility for the legitimate intelligence of ethical judgments that are not conceptual in nature.

---

[219] See J.J. Thomson, "The Trolley Problem," *Yale Law Journal*, 94 (1985): 1395-1415; Joshua Greene, "Why Are VMPFC Patients More Utilitarian?: A Dual-Process Theory of Moral Judgment Explains," *Trends in Cognitive Sciences* 11(2007): 322-323 for issues with giving rational rule application explanations of ethical judgments.

[220] Quite obviously, we should think of Aristotle, The *Nicomachean Ethics* (Oxford: Oxford University Press, 1988) and habit here. But the intelligence of virtue ethics, on my account, can be explained as something more than conditioned responses. On a skill-based account of ethics, the intellect of ethical judgments is not reduced to brute conditioning or propositional knowledge.

This same notion can be applied to aesthetic judgments. After all, the sensitivity to color, light, proportion, harmony etc., involved in aesthetic judgments, like that of skills, seems to defy explanation in terms of traditional truth-preserving logical reasoning. However, it is hardly adequate to describe aesthetic judgments as brute or mechanistic. The possibility of approaching aesthetic judgments in a way that is similar to the practical intelligence of skill offers some semblance of hope for presenting a thorough explication of the type of cognition that is involved in aesthetic reasoning. Overall, my proposal for the diachronic cognitive penetration of perception by skill opens up several avenues for thinking about domains of thought that defy analysis in terms of traditional rationality, but remain squarely cognitive.

3. Future Research

3.1 Cognition

As I stated earlier, assessing whether or not perception is cognitively penetrable will depend crucially on determining which states and processes qualify as cognitive. In fact, the issues above that arise in connection with sensorimotor skills are, at heart, issues about the limits of cognition. After all, the only way that we will be able to circumscribe, in a fully transparent fashion, the limits of intentional action (Criterion (1)) and practice (Criterion (3)) will be to investigate the nature and scope of cognition.

In this dissertation, as far as possible, I have tried to avoid the peripheral edges of cognition. I have intentionally focused on behaviors and abilities that are unobjectionably intelligent. Inevitably, however, if we are to decide on limit cases, we must answer the question "what is cognition?" This is imperative not only for dealing with issues of perception, but for securing the theoretical foundations of a cognitive science, generally.

Though I have not directly considered the nature of cognition in this dissertation, I do think that I have illuminated some relevant considerations for developing a definition of this term. Importantly, I have highlighted the fact that a definition of cognition must do more than give an account of propositional thought and reasoning.[221] The data that must be explained by a theory of cognition must include processes, behaviors and non-linguistic, non-propositional representational states. If a definition of cognition does not incorporate these types of states and events under it, then its usefulness for cognitive science will be severely limited.

In thinking about circumscribing a definition of cognition, it should be clear that traditional notions of rationality as the ability to do logical deduction will not serve us well in this context. We can see this quite clearly if we imagine a very simple computing system that is programmed to deduce from premise (1) A v B and premise (2) –B the conclusion that A. The conclusion of course is perfectly valid, but it would be a stretch to call this system cognitive or intelligent simply in virtue of its capacity to deduce a valid conclusion. Quite clearly then, it follows that reasoning validly from a set of premises to a conclusion is not a sufficient condition for intelligence. Otherwise, any system that could do deduction of the sort described above would be considered intelligent.

The important question, however, is why does this system lack intelligence? I claim that what is highlighted by our unwillingness to bestow upon the disjunction machine the quality of intelligence is the fact that what makes a state, process or behavior intelligent is not that it is an instance of logical deduction, but rather that it is characterized by properties of some other sort. The important question is, of course, "which properties?"

[221] This is not to say that an account of propositional thought is unnecessary or unimpressive. It is to say, however, that an account of propositional thought is not sufficient as an account of cognition. And an account of cognition is exactly what concerns me here.

It cannot be the possession of intentional states because, as I've argued, we want to remain neutral on whether non-linguistic animals and pre-linguistic children can possess cognitive states. But, if propositional states are the way by which we pick out cognition, then we would either have to exclude such creatures from participating in the cognitive realm or we will have to let them in by reasoning in a circle. As I've argued above, making representational states a prerequisite for cognition distracts us from the characteristics that we are actually using to pick out cognitive behavior. This is because the way we decide if non-linguistic creatures possess representational states is by observing their behavior and determining if it is sophisticated enough to warrant the attribution of representational states. But it is supposed to be the presence of representational states that determines whether or not a behavior is cognitive. Such a definition of cognition will simply not do.

Additionally, the activity of particular anatomical brain structures cannot be the way by which we determine cognitive states or behaviors either.[222] This is the case because such a definition of cognition necessarily excludes creatures that don't have brains like ours from participation in the cognitive realm. For instance, if cortical activation is necessary for possession of cognitive states or the expression of cognitive behavior, then creatures such as octopi and computers will necessarily be noncognitive. But such a result is quite strange—there seems to be no principled reasons why octopi could not possess cognitive states.

I am convinced that a satisfactory definition of cognition will isolate and elucidate those properties that we regularly use to determine the intelligence of states and behaviors. Defining cognition by way of characteristic properties, rather than by identifying the types of things that may possess these properties, I believe, will allow us to produce the most useful and effective definition of cognition. That is, a satisfactory definition of cognition should, as far as possible,

[222] See Jesse Prinz, *Gut Reactions: A Perceptual Theory of Emotion* (Oxford: Oxford University Press, 2004), 45-49.

stay neutral on the ontological status of those things that may be characterized by cognition, it should allow us to decide, in its application, whether a process, event, state or behavior qualifies as cognitive. We should be able to do this not by way of identifying the sort of thing that it is, but by determining the kinds of properties that it has. Any other way of identifying cognition, I fear, will simply expose our prejudices against certain kinds of states, events or creatures.

I think that the best methodology for isolating these features will be to pick out paradigmatically cognitive states and events, identify their central features and then subtract incompatible properties. So, we may look at reasoning and problem solving as paradigmatically cognitive activities, but since reasoning is necessarily over propositional states but problem solving is not, we can subtract propositionality from the necessary conditions of cognition. But, what remains?

I believe that the following three components will be crucial to any successful definition of cognition: (1) any cognitive state or behavior will be goal-directed or goal-oriented.[223] (2) Any cognitive state or behavior must be sensitive to variation in its surroundings (environmental or conceptual). And (3) any cognitive state or behavior should possess the capacity for revision. Defending this notion and providing the relevant thresholds will constitute the next step of my research.

2.2 Enactive Emotion

It is a fairly common belief—the sort of thing that new-agey people often repeat, that experience is a matter of perspective. The origin of this watered-down wisdom is found in Buddhist philosophy. A central tenet in Buddhist philosophy is the idea that one can overcome the suffering caused by a world in flux, by changing what one desires. One can overcome suffering (snuff it out, not just ignore it), by overcoming attachment. The reasoning is at very

[223] See Robert Nozick, *The Nature of Rationality* (Princeton, NJ: Princeton University Press, 1993), 133.

core of Buddhist thought and it is very simple; it is the four noble truths: (1) life is suffering, (2) suffering has a cause, (3) the cause of suffering can be overcome, (4) the cause of suffering can be overcome by following the eight-fold path.[224]

First off, it is important to notice that suffering is a state of a person, and not a state of the world;[225] it is an emotional state and it is the result of loss. Since even the most permanent-seeming things pass, suffering is central to human existence. That is, the fact of impermanence ensures suffering. But, impermanence is not the cause of suffering, it is simply the condition that makes suffering inevitable if its cause is not removed. The actual cause of suffering is desire. It is attachment to the objects of this world that causes suffering because one always mourns the loss of those things to which one is attached. But there is no way not to lose things because impermanence is a categorical condition of the world. So, if one is attached to things, one will inevitably suffer their loss. Optimistically, the cause of suffering can be overcome by extinguishing desire and attachment. By not having objects of attachment, one need not mourn their loss, and thus, one can avoid suffering. By following a set of guidelines, one can change what one desires, which will in turn alleviate suffering. Importantly, how one lives and what one believes and desires can change how one feels. That is, by changing our desires we can eliminate feelings of suffering.

Now, you may be asking, "how this is relevant?" It is relevant because central to this Buddhist doctrine is the idea that we can change, not just the way we *respond* to our hardwired feelings, but the actual *constitution* of those feelings. That is, what one believes and desires, one's intentional states, influence what one feels—the qualitative character of one's emotions. In Buddhism, one does not just learn to ignore suffering but one changes what one feels by

[224] See Walpola Rahula, *What the Buddha Taught* (New York: Gove Press, 1959), 16-28.
[225] Rahula, 31.

thinking and living in a particular fashion. One does not simply learn to withstand suffering;[226] one learns to extinguish it.

Compare this to Aristotle's position concerning the emotions. According to Aristotle, the passions are given to us by nature, [227] and "nothing that exists by nature can form a habit contrary to its nature."[228] Quite clearly then, this means that one does not have the power to changes one's emotions. One can control how one reacts to those emotions—one can fight or suppress her natural inclinations, but one cannot change the actual constitution of those inclinations. If we naturally feel anger when slighted, we will always feel anger when slighted. Beliefs and desires, on an Aristotelian model, do not impact feeling. [229] Another way of putting this is: the processing of the qualitative character of an emotional state is informationally encapsulated.

The reason that I bring this up here is because these opposing positions regarding the changeability of the emotions are strikingly similar to the debate about perception and its relationship to cognition. Moreover, not only are the debates surrounding emotions and perceptions similar, but descriptions of these states are striking in their similarity as well. That is, both emotions and perceptions, unlike all other mental states, are described as hybrid states. Hybrid states are constituted by a qualitative and an intentional component. And just like with perceptions, it is only the qualitative component's response to cognition that is in any way

---

[226] Buddhists are not Stoics.

[227] "[W]e have faculties by nature, but we are not made good or bad by nature" Aristotle, 1106a8; "Again, we feel anger and fear without choice, but the virtues are modes of choice or involve choice" Aristotle, 1106a1; "[W]e are not called good or bad on the ground of our passions, but are so called on the ground of our virtues and our vices, and because we are neither praised nor blamed for our passions (for the man who feels fear or anger is not praised, nor is the man who simply feels anger blamed, but the man who feels it in a certain way), but for our virtues and our vices we *are* praised and blamed" Aristotle, 1105b29-34.

[228] Aristotle, 1103a14.

[229] Aristotle is not saying that judgment does not effect emotion at all, since clearly the appropriateness of a response will depend on the judgment of the nature of a situation. If one does not believe that he is being wronged, he will not feel angry. In this way, the assessment of the type of situation that one is encountering is quite uncontroversially able to impact how one feels about the situation. The deeper claim that I am exploring here, however, is the idea that after one judges the situation to be of a particular kind, whether one's beliefs, desires and knowledge can affect how one will feel in response to that situation.

controversial. Everyone agrees that the intentional aspect of both perceptions and emotions are cognitively penetrated.

Because of these similarities, the potential significance of my dissertation for a theory of emotion is enormous. Specifically, I am proposing that my arguments for the diachronic cognitive penetration of perception by skill may be applicable to emotions. Thus, if it is possible that skills impact the qualitative character of a perceptual event, there is also the distinct possibility that emotions may be impacted by skills in this way, as well.

REFERENCES

Anscombe, Elizabeth. *Intention.* Cambridge, MA: Harvard University Press, 1957.

Aristotle. *The Nicomachean Ethics.* Oxford: Oxford University Press, 1988.

Aristotle. "On the Soul." In *The Complete Works of Aristotle: Volume One,* edited by Jonathan Barnes. Princeton: Princeton University Press, 1984.

Armstrong, David. *A Materialist Theory of the Mind.* London: Routledge and Kegan Paul, 1968.

Austin, J.L. "A Plea for Excuses." In *Classics of Analytic Philosophy,* edited by Robert Ammerman, 379-398. USA: Hacket Publishing, 1990.

Bao, Shaowen, Chang, Edward F., Woods, Jennifer and Merzenich, Michael M. "Temporal Plasticity in the Primary Auditory Cortex Induced by Operant Perceptual Learning." *Nature Neuroscience* 7 (2004): 974-981.

Bengson, John and Moffet, Marc. "Know-How and Concept-Possession." *Philosophical Studies* 136 (2007): 387-401.

Bengson, John, Moffet, Marc and Wright, Jennifer. "The Folk on Knowing-How." *Philosophical Studies* 142 (2009): 387-401.

Berkeley, George. *A Treatise Concerning the Principles of Human Knowledge.* Oxford: Oxford University Press, 1998.

Bermúdez, José. "Animal Reasoning and Proto-Logic." In *Rational Animals?* edited by Susan Hurley and Matthew Nudds, 127-138. Oxford: Oxford University Press, 2006.

Bermúdez, José. *Thinking without Words*. Oxford: Oxford University Press, 2003.

Berry, D.C. and Broadbent, D.E. "On the Relationship Between Task Performance and Associated Verbalizable Knowledge." *Quarterly Journal of Experimental Psychology* 36 (1984): 209-231.

Black, J.E., Isaacs, K.R., Anderson, B.J., Alcantara, A.A. and Greenough, W.T. "Learning Causes Synaptogenesis, Whereas Motor Activity Cause Angiogenesis in Cerebellar Cortex of Adult Rats." *Proceedings of National Academy of Sciences MSA* 87 (1990): 55568-5572.

Borat. Directed by Larry Charles. New York: 20th Century Fox, 2006.

Boyson, S.T. and Bernston, G. "Responses to Quantity: Perceptual Versus Cognitive Mechanisms in Chimpanzees (*Pan troglodytes).*" *Journal of Experimental Psychology and Animal Behavior Processes* 21 (1995): 82-86.

Bransford, John D., Brown, Ann L. and Cocking, Rodney R. *How People Learn: Brain, Mind, Experience and School.* Washington D.C.: National Academy Press, 1999.

Bratman, Michael. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press, 1987.

Brehmer, B., Hagafors, R. and Johansson, R. "Cognitive Skills in Judgment: Subjects' Ability to Use Information about Weights, Function Forms, and Organizing Principles." *Organizational Behavior and Human Performance* 26 (1980): 373-85.

Brentano, Franz. *Psychology from an Empirical Standpoint*. London: Routledge, 1973.

Broadbent, D.E. and Aston, B. "Human Control of a Simulated Economic System." *Ergonomics* 21 (1978), 1035-1043.

Bruner, Jerome S. and Potter, Mary C. "Interference in Visual Recognition." *Science* 144 (1964): 424-5.

Bruner, Jerome S. and Postman, Leo. "On the Perception of Incongruity: A Paradigm." *Journal of Personality* 18 (1949): 206-23.

Bruner, Jerome S. and Mintum, Ann Leigh. "On Perceptual Readiness." *Psychological Review* 64 (1957): 123-52.

Bruner, Jerome S. and Mintum, Ann Leigh. "Perceptual Identification and Perceptual Organization." *Journal of General Psychology* 53 (1955): 21-28.

Bullier, Jean. "Visual Perception is Too Fast to be Impenetrable to Cognition." In "Open Peer Commentary: Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," by Zenon Pylyshyn *Behavioral and Brain Sciences* 22 (1999): 370.

Byrne, Alex. "Intentionalism Defended." *Philosophical Review* 110 (2001): 199-239.

Campbell, John. "Manipulating Colour: Pounding an Almond." In *Perceptual Experience,* edited by Tamar Szabo Gendler and John Hawthorne, 31-48. Oxford: Oxford University Press, 2006.

Campbell, John. *Past, Space and Self.* Cambridge, MA: The MIT Press, 1995.

Campbell, John "Transparency vs. Revelation in Color Perception." *Philosophical Topics* 33, (2005): 105-115.

Carlson, Richard A. "Skill learning." In *The Encyclopedia of Cognitive Science,* edited by L. Nadal. London: Macmillan Reference Ltd, 2003.

Carr, David. "The Logic of Knowing How and Ability." *Mind* 88 (1979): 394-409.

Carr, David. "Knowledge in Practice." *American Philosophical Quarterly* 8 (1981): 53-61.

Carvallo, R.P., Tudella, E., Cljouw, S.R. and Savelsbergh, G.J. "Early Control of Reaching: Effects of Experience and Body Orientation." *Infant Behavioral Development* 31 (2008): 23-33.

Chase, W.G. and Simon, H.A. "The Mind's Eye in Chess." In *Visual Information Processing,* edited by W. G. Chase, 215-81. New York: Academic Press, 1973.

Cheng, P.W. "Restructuring Versus Automaticity: Alternate Accounts of Skill Acquisition." *Psychological Review* 92 (1985): 414-423.

Chomsky, Noam. *Aspects of the Theory of Syntax.* Cambridge, MA: The MIT Press, 1965.

Cheung, Steven W., Nagarajan, Srikantan S., Schreiner, Christoph E., Bedenbaugh, Purvis H. and Wong, Andrew. "Plasticity in Primary Auditory Cortex of Monkeys with Altered Vocal Production." *The Journal of Neuroscience* 25 (2005): 2490-2503.

Churchland, Paul and Churchland, Patricia. "Could A Machine Think?" *Scientific American* 262 (1990): 32-37.

Churchland, Paul. "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." *Philosophy of Science* 55 (1988): 167-87.

Clark, Andy and Eilan, Naomi. "Sensorimotor Skills and Perception." *Aristotelian Society Supplementary* 80 (2007): 67-88.

Clark, Andy. *Supersizing the Mind.* Oxford: Oxford University Press, 2008.

Clark, Austen. "Feature-Placing and Proto-Objects." *Philosophical Psychology* 17 (2004): 443-469.

Clark, Austen. *Sensory Qualities.* Oxford: Oxford University Press, 1993.

Cole, Jonathan and Paillard, Jacques. "Living Without Touch and Peripheral Information About Body Position and Movement: Studies with Deafferented Subjects." In *The Body and the Self,* edited by José Bermúdez, Anthony Marcel and Naomi Eilan, 245-66. Cambridge, MA: The MIT Press, 1995.

Crane, Tim. *The Contents of Experience: Essays on Perception.* Cambridge: Cambridge University Press, 1992.

Crane, Tim. "The Waterfall Illusion." *Analysis* 48 (1988): 142-147.

Cummins, Robert. *Meaning and Mental Representation.* Cambridge, MA: The MIT Press, 1989.

Cussins, Adrian. "The Connectionist Construction of Concepts." In *The Philosophy of Artificial Intelligence,* edited by Margaret A. Boden, 368-440. Oxford: Oxford University Press, 1990.

Cussins, Adrian. "Content, Embodiment and Objectivity: The Theory of Cognitive Trails." *Mind*
101 (1992): 651-88.

Davidson, Donald. *Essays on Actions and Events.* Oxford: Clarendon Press, 2001.

Davidson, Donald. "Knowing One's Own Mind." *Proceedings and Addresses of the American Philosophical Association* 60 (1987): 441-458.

Davidson, Donald. "Thought and Talk." In *Mind and Language,* edited by Samuel Guttenplan, 7-23. Oxford: Oxford University Press, 1975.

Davies, Ian R.L., Sowden, Paul T., Hammond, Sean M. and Ansell, Janet. "Expertise in Categorizing Mammograms: A Perceptual or Conceptual Skill?" In *Medical Imaging 1994: Image perception,* edited by Harold L. Kundel. *Proceedings of SPIE 2166* (1994): *86-94.*

Dennett, Daniel. "Can Machines Think?" In *How We Know,* edited by Michael Shafto 121-45. San Francisco, CA: Harper and Row, 1984.

Dennett, Daniel. *Content and Consciousness.* London: Routledge and Kegan Paul, 1969.

Dennett, Daniel. *The Intentional Stance.* Cambridge, MA: The MIT Press 1989.

Dennett, Daniel. "Ways of Establishing Harmony. " In *Dretske and His Critics,* edited by Brian P. McLaughlin, 118-130. Cambridge, MA: Blackwell Publishing, 1991.

Descartes, Rene. *Meditations on First Philosophy.* Indianapolis, IN: Hacket, 1993.

Devitt, Michael. *Ignorance of Language.* Oxford: Oxford University Press, 2006.

Devitt, Michael. "Methodology and the Nature of Knowing How." Forthcoming.

Doerfler, R.W. *Dead Reckoning: Calculating without Instruments.* Houston: Gulf Publishing, 1993.

Dretske, Fred. *Explaining Behavior.* Cambridge, MA: The MIT Press, 1988.

Dretske, Fred. *Knowledge and the Flow of Information.* Cambridge, MA: The MIT Press, 1981.

Dretske, Fred. "Machines and the Mental." *Proceedings and Addresses of the American Philosophical Association* 59 (1985): 23-33.

Dretske, Fred. "Minimal Rationality." In *Rational Animals?* edited by Susan Hurley and Mathew Nudds, 107-8. Oxford: Oxford University Press, 2006.

Drestke, Fred. *Naturalizing the Mind.* Cambridge, MA: The MIT Press, 1997.

Dretske, Fred. *Perception, Knowledge and Belief.* Cambridge: Cambridge University Press, 2000.

Dretske, Fred. "Seeing, Believing and Knowing." In *Perception,* edited by Robert Schwartz, 268-86. Oxford: Blackwell Publishing, 2004.

Dreyfus, Hubert L. "The Return of the Myth of the Mental." *Inquiry* 50 (2007): 352-65.

Dummett, Michael. *Origins of Analytic Philosophy*. Cambridge, MA: Harvard University Press, 1993.

Ericsson, K. Anders and Charness, Neil. "Expert Performance: Its Structure and Acquisition." *American Psychologist* 49 (1994): 725-74.

Evans, Gareth. *The Varieties of Reference.* Oxford: Oxford University Press, 1982.

Fahle, Manfred. "Perceptual Learning: A Case for Early Selection." *Journal of Vision* 4 (2004): 879-90.

Fodor, Jerry. " A Reply to Churchland's 'Perceptual Plasticity and Theoretical Neutrality'." *Philosophy of Science* 55 (1988): 188-98.

Fodor, Jerry. *A Theory of Content and Other Essays*, Cambridge, MA: The MIT Press, 1990.

Fodor, Jerry and Pylyshyn, Zenon. "How Direct is Visual Perception?: Some Reflections on Gibson's "Ecological Approach." In *Vision and Mind: Selected Readings in the Philosophy of Perception*, edited by Alva Noë & Evan Thompson, 167-228. Cambridge, MA: The MIT Press, 2002.

Fodor, Jerry. *The Language of Thought*, Cambridge, MA: Harvard University Press, 1975.

Fodor, Jerry. *LOT 2: The Language of Thought Revisited*. Oxford: Oxford University Press, 2008.

Fodor, Jerry. *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. Cambridge, MA: The MIT Press, 2000.

Fodor, Jerry. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: The MIT Press, 1983.

Fodor, Jerry. "Modules, Frames, Fridgeons, Sleeping Dogs and the Music of Spheres." In *The Robot's Dilemma: The Frame Problem in Artificial Intelligence,* edited by Zenon Pylyshyn, 139-49. Norwood, NJ: Ablex Publishing, 1987.

Fodor, Jerry. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: The MIT Press, 1987.

Frankfurt, Harry. "The Problem of Action." *American Philosophical Quarterly* 15 (1978): 157-62.

Fritz, Jonathan, Shamma, Shihab, Elhilali, Mounya and Klein, David. "Rapid Task-Related Plasticity of Spectrotemporal Receptive Fields in Primary Auditory Cortex." *Nature Neuroscience* 6 (2003):1216-223.

Geach, Peter. "Identity." *Review of Metaphysics* 21 (1967): 3-12.

Gibson, J.J. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin, 1979.

Ginet, Carl. *Knowledge, Perception and Memory.* The Netherlands: D. Reidel Publishing Company, 1975.

Goodale, M.A. "Modularity in the Visuomotor Control: From Input to Output." In *Computational Processes in Human Vision: An Interdisciplinary Perspective,* edited by Zenon Pylyshyn, 262-85. Norwood, NJ: Ablex Publishing, 1988.

Goodale, M.A. "Vision as a Sensorimotor System." In *Behavioral Approaches to Brain Research,* edited by T.E. Robinson, 41-61. Oxford: Oxford University Press, 1983.

Gonzalez-Lima, F., Ferchmin, P. A., Eterovic, V. A. and Gonzalez-Lima, E. M. "Metabolic Activation of the Brain of Young Rats After Exposure to Environmental Complexity." *Developmental Psychobiology* 27 (1993): 343-51.

Gopnik, Alison. "Your Baby is Smarter Than You Think." New York Times, August

15, 2009, Opinion, New York edition.

Greene, Joshua. "Why are VMPFC Patients More Utilitarian?: A Dual-Process Theory of Moral Judgment Explains." *Trends in Cognitive Sciences* 11(2007): 322-323.

Grunewald, Alexander. "Neurophysiology Indicates Cognitive Penetration of the Visual System." In "Open Peer Commentary: Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception," by Zenon Pylyshyn *Behavioral and Brain Sciences* 22 (1999): 379-80.

Hardcastle, Valerie Gray. "Computationalism." *Synthese* 105 (1995): 303-17.

Harman, Gilbert. "Practical Reasoning." In *The Philosophy of Action*, edited by Alfred R. Mele. Oxford: Oxford University Press, 1997.

Hawley, Katherine. "Success and Knowledge How." *American Philosophical Quarterly* 40 (2003): 19-31.

Hetherington, Stephen. "How to Know (That Knowledge-That is Knowledge-How)." In *Epistemology Futures*, edited by Stephen Hetherington, 71-94. Oxford: Oxford University Press, 2006.

Held, Richard and Hein, Alan. "Movement-Produced Stimulation in the Development of Visually Guided Behavior." *Journal of Comparative and Physiological Psychology* 56 (1963): 872-876.

Hume, David. *An Enquiry Concerning Human Understanding.* Indianapolis, IN: Bobbs-Merrill, 1955.

Hume, David. *Treatise of Human Nature.* USA: Prometheus Books, 1992.

Hurley, Susan. *Consciousness in Action.* Cambridge, MA: Harvard University Press, 1998.

Hurley, Susan. "Making Sense of Animals." In *Rational Animals?* edited by Susan Hurley and Matthew Nudds, 139-171. Oxford: Oxford University Press, 2006.

Jackson, Frank. "What Mary Didn't Know." *Journal of Philosophy* 83 (1986): 291-5.

Kacelnik, Alex. "Meanings of Rationality." In *Rational Animals?* edited by Susan Hurley and Matthew Nudds, 87-106. Oxford: Oxford University Press, 2006.

Kelly, Sean. "Demonstrative Concepts and Experience." *The Philosophical Review* 110 (2001): 397-420.

Koethe, John. "Comments and Criticism: Stanley and Williamson on Knowing How."

Journal of Philosophy 99 (2002): 325-8.

Kripke, Saul. *Naming and Necessity*. Cambridge, MA: Harvard University Press, 1972.

LaBerge, D. and Samuels, S.J. "Toward a Theory of Automatic Information Processing in Reading." *Cognitive Psychology* 6 (1974): 293-323.

Lewicki, P., Hill, T. and Bizot, E. "Acquisition of Procedural Knowledge about a Pattern of Stimuli that Cannot be Articulated." *Cognitive Psychology* 20 (1988): 24-37.

Li, Wu, Piëch, Valentin and Gilbert, Charles D. "Perceptual Learning and Top-Down Influences in Primary Visual Cortex." *Nature Neuroscience* 7 (2004): 651-57.

Locke, John. *An Essay on Human Understanding.* London: Penguin Books, 1997.

Loewer, Barry. "From Information to Intentionality." *Synthese* 70 (1987): 287-317.

Logan, G.D. "Automaticity, Resources, and Memory: Theoretical Controversies and Practical Implications." *Human Factors* 30 (1988): 583–598.

Lycan, William. "In Defense of the Representational Theory of Qualia (Replies to Neander, Rey and Tye)." In *Philosophical Perspectives* 12 (1998): 479-87.

Maunsell, J.H.R., Sclar, G., Nealey, T.A. and Depriest, D.D. "Extraretinal Representations in Area V4 in the Macaque Monkey." *Visual Neurosceince* 7(1991): 561-73.

McIntosh, A.R. "Network Analysis of Cortical Visual Pathways Mapped with PET." *Journal of Neuroscience,* 14 (1994): 655-66.

McDowell, John. *Mind and World*. Cambridge, MA: Harvard University Press, 1994.

Mele, Alfred R. and Moser, Paul K. "Intentional Action." *Noûs* 28 (1994): 39-68.

Merleau-Ponty, Maurice. *The Phenomenology of Perception.* London: Routledge, 1945.

Millar, Alan. *Reasons and Experience.* Oxford: Clarendon Press, 1991.

Millikan, Ruth Garrett. "Styles of Rationality." In *Rational Animals?* edited by Susan Hurley and Mathew Nudds, 117-126. Oxford: Oxford University Press, 2006.

Miyashita, Yasushi and Hayashi, Toshihiro. "Neural Representation of Visual Objects: Encoding and Top-down Activation." *Current Opinion in Neurobiology*, 10 (2000): 187-194.

Moore, G.E. "A Reply to My Critics." In *The Philosophy of G. E. Moore,*

edited by P. A. Schilpp, 535-677. Evanston, IL: Northwestern University Press, 1942.

Newell, K.M. "Motor Skill Acquisition." *Annual Review of Psychology* 42 (1991): 213-37.

Noë, Alva. *Action in Perception.* Cambridge, MA: The MIT Press, 2004.

Noë, Alva. "Against Intellectualism." *Analysis* 65 (2005): 278-290.

Noë, Alva. "Experience without the Head." In *Perceptual Experience,* edited by Tamar Szabo Gendler and John Hawthorne, 411-433. Oxford: Oxford University Press, 2006.

Noë, Alva and O'Regan, Kevin. "On the Brain-Basis of Visual Consciousness: A Sensori-Motor Account." In *Vision and Mind: Selected Readings in the Philosophy of Perception,* edited by Alva Noë & Evan Thompson, 567-98. Cambridge, MA: The MIT Press, 2002.

Nozick, Robert. *The Nature of Rationality.* Princeton, NJ: Princeton University Press, 1993.

O'Toole, A.J. and Kersten, D.J. "Learning to See Random-Dot Stereograms." *Perception* 21(1992): 227-43.

Palmer, S.E. "The Effects of Contextual Scenes on the Identification of Objects." *Memory and Cognition* 3 (1975): 519-26.

Pantev, Christo, Ross, Bernhard, Fujioka, Takako, Trainor, Laurel J., Schulte, Michael and Schulz, Mitthias. "Music and Learning-Induced Cortical Plasticity." In *The Neurosciences and Music,* edited by Giuliano Avanzini, Carmine Faienza, Diego Minciacchi, Luisa Lopez and Maria Majno, 438-450. New York: New York Academy of Sciences, 2003.

Papineau, David and Heyes, Cecelia. "Rational or Associative? Imitation in Japanese Quail." In *Rational Animals?* edited by Susan Hurley and Mathew Nudds, 187-195. Oxford: Oxford University Press, 2006.

Peacocke, Christopher. "Analogue Content." *Proceedings of the Aristotelian Society* 60 (1986): 1-17.

Peacocke, Christopher. *A Study of Concepts.* Cambridge: MIT Press, 1992.

Peacocke, Christopher. "Perceptual Content." In *Themes from Kaplan,* edited by Joseph Almog, John Perry, and Howard Wettstein, 297-329. New York: Oxford University Press, 1989.

Pearson, Richard G., Hauty, George T. "Adaptive Processes Determining Proprioceptive

Perception of Verticality." *Journal of Experimental Psychology* 57 (1959): 367-371.

Pourtois, Gilles, Rauss, Karsten S., Vuilleumier, Patrik and Schwartz, Sophie. " Effects of Perceptual Learning on Primary Visual Cortex Activity in Humans." *Vision Research* 48 (2008): 55-62l.

Prinz, Jesse. "Beyond Appearances: The Content of Sensation and Perception." In *Perceptual Experience,* edited by Tamar Szabo Gendler and John Hawthorne, 434-460. Oxford: Oxford University Press, 2006.

Prinz, Jesse. *Gut Reactions: A Perceptual Theory of Emotion.* Oxford: Oxford University Press, 2004.

Prinz, Jesse. "Is the Mind Really Modular?" In *Contemporary Debates in Cognitive Science,* edited by Robert J. Stainton, 22-36. Malden, MA: Blackwell Publishing, 2006.

Pylyshyn, Zenon. "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception." *The Behavioral and Brain Sciences* 22 (1999): 343-91.

Pylyshyn, Zenon. *Seeing and Visualizing.* Cambridge, MA: MIT Press, 2003.

Raftopoulos, Athanassio. "Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception." *Cognitive Science* 25 (2001): 423-451.

Rahula, Walpola. *What the Buddha Taught.* New York: Gove Press, 1959.

Reber, Arthur S. and Millward, R.B. "Event Tracking in Probability Learning." *American Journal of Psychology* 84 (1971): 85-99.

Reber, Arthur and Lewis, Selma. "Implicit Learning: An analysis of the Form and Structure of a Body of Tacit Knowledge." *Cognition* 5 (1977): 333-61.

Reber, Arthur S. "Implicit Learning and Tacit Knowledge*." Journal of Experimental Psychology: General* 118(1989), 219-35.

Recanzone, G.H., Schreiner, C.E. and Merzenich, M.M. "Plasticity in the Frequency Representation of Primary Auditory Cortex Following Discrimination Training in Adult Owl Monkeys." *The Journal of Neuroscience* 13 (1993): 87-103.

Rosenbaum, David A., Carlson, Richard A., Gilmore, Rick O. "Acquisition of Intellectual and Perceptual Motor Skills." *Annual Review of Psychology* 52 (2001): 453-70.

Rosenthal, David M. "Being Conscious of Ourselves." *The Monist* 87 (2004): 161-84.

Rosenthal, David M. *Consciousness and Mind.* Oxford: Oxford University Press, 2005.

Rosenthal, David M. "State Consciousness and Transitive Consciousness." *Consciousness and Cognition* 2 (1994): 355-63.

Rosenthal, David M. "Two Concepts of Consciousness." In *The Nature of Mind*, edited by David M. Rosenthal, 462-77. New York: Oxford University Press, 1991.

Rosenzweig, Mark R. and Bennett, Edward L. "Cerebral Changes in Rats Exposed Individually to an Enriched Environment." *Journal of Comparative and Physiological Psychology* 8 (1972): 304-313.

Ryle, Gilbert. *The Concept of Mind.* Chicago: The University of Chicago Press, 1949.

Sacks, Oliver. *An Anthropologist on Mars: Seven Paradoxical Tales.* New York: Knopf, 1995.

Sacks, Oliver. *The Man Who Mistook His Wife For A Hat And Other Clinical Tales.* New York: Touchstone Books, 1985.

Saling, L. and Philips, J.G. "Automatic Behavior: Efficient Not Mindless." *Brain Research Bulletin* 73 (2007): 1–20.

Sanes, Jerome N. and Donogue, John P. "Plasticity and Primary Motor Cortex." *Annual Review of Neuroscience* 23 (2000): 393-415.

Schui, L. and Pashler, H. "Improvement in Line Orientation Discrimination is Retinally Local but Dependent on Cognitive Set." *Perception and Psychophysics* 52(1992): 582-88.

Schwartz, Robert. *Perception: Blackwell Readings in Philosophy.* Malden, MA: Blackwell Publishing, 2004.

Searle, John. "Minds, Brains and Programs." *Behavioral and Brain Sciences* 3 (1980): 417-57.

Siegal, Susanna. "Which Properties are Represented in Perception?" In *Perceptual Experience*, edited by Tamar Szabo Gendler and John Hawthorne, 481-503. Oxford: Oxford University Press, 2006.

Seitz, Aaron R. and Dinse, Hubert R. "A Common Framework for Perceptual Learning." *Current Opinion in Neurobiology* 17 (2007): 1-6.

Sellars, Wilfrid. *Empiricism and the Philosophy of Mind.* Cambridge, MA: Harvard University Press, 1997.

Sellars, Wilfrid. "Philosophy and the Scientific Image of Man." In *Frontiers of Science and Philosophy*, edited by Robert Colodny, 35-78. Pittsburgh, PA: University of Pittsburgh Press, 1962.

Smith, S.B. *The Great Mental Calculators: The Psychology, Methods, and Lives of Calculating Prodigies, Past and Present*. New York: Columbia University Press, 1983.

Snowdon, Paul. "Knowing How and Knowing That: A Distinction Reconsidered." *Proceedings of the Aristotelian Society* 104 (2003): 1–29.

Sowden, P.T., Davies, I.R.L., Rose, D. and Kaye, M. "Perceptual Learning of Stereoacuity." *Perception* 25 (1996): 1043-52.

Spelke, Elizabeth S. "Where Perceiving Ends and Thinking Begins: the Apprehension of Objects in Infancy." In *Perception*, edited by Robert Schwartz, 321-36. Oxford: Blackwell Publishers, 2004.

Stanley, Jason. "Knowing (How)." Forthcoming in *Nous.*

Stanley, Jason and Williamson, Timothy. "Knowing How." *Journal of Philosophy* 98 (2001): 411-44.

Staszewski, J.J. "Skilled Memory and Expert Mental Calculation." In *The Nature of Expertise,* edited by M.T.H. Chi, R. Glaser and M.J. Farr, 71-128. Hillsdale, NJ: Erlbaum, 1988.

Strawson, P. F. *Individuals: An Essay in Descriptive Metaphysics*. New York: Anchor Books, 1959.

Strawson, P.F. "Perception and Its Objects." In *Vision and Mind: Selected Reading in the Philosophy of Perception,* edited by Alva Noë and Evan Thompson, 91-110. Cambridge, MA: The MIT Press, 2002.

Strawson, P.F. *Subject and Predicate in Logic and Grammar.* London: Methuen, 1974.

Thomson, J.J. "The Trolley Problem." *Yale Law Journal*, 94 (1985): 1395-1415.

Turing, Alan. "Computing Machinery and Intelligence." *Mind* 59 (1950): 433-60.

Tye, Michael. "Nonconceptual Content, Richness, and Fineness of Grain." In *Perceptual Experience,* edited by Tamar Szabo Gendler and John Hawthorne, 504-530. Oxford: Oxford University Press, 1996.

Tye, Michael. *Ten Problems of Consciousness*. Cambridge: The MIT Press, 1995.

Wallace, David Foster. "Federer as Religious Experience." New York Times, August 20, 2006, Play Magazine, New York edition.

Wallis, Charles. "Consciousness, Context and Know-how." *Synthese* 160 (2008): 123-53.

Williams, John N. "Propositional Knowledge and Know-How." *Synthese* 165 (2008): 107-25.

Wittgenstein, Ludwig. *Philosophical Investigations*. Oxford: Blackwell Publishing, 1953.

Wright, Crispin. "Wittgenstein's Later Philosophy of Mind: Sensation, Privacy, and Intention." *Journal of Philosophy* 86 (1989): 622-34.