Summer 8-15-2021

The Disunity of Perception

Benjamin Austin Henke
Washington University in St. Louis

WASHINGTON UNIVERSITY IN ST. LOUIS

Department of Philosophy
Program in Philosophy-Neuroscience-Psychology

Dissertation Examination Committee:
Casey O'Callaghan, Chair
Allan Hazlett
Ron Mallon
Jake Quilty-Dunn
Julia Staffel
Daniel Stoljar

The Disunity of Perception

by

Benjamin Henke

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

August 2021
St. Louis, Missouri

Table of Contents

List of Figures

Acknowledgements

I’m immensely grateful to my wife, Julia Haas, for her constant support throughout this process and at every stage of writing, editing, depression, writing again, and submission. I simply could not have done this without her. I’m also grateful to my committee — Casey O’Callaghan, Allan Hazlett, Ron Mallon, Jake Quilty-Dunn, Julia Staffel, and Daniel Stoljar — for their immense help along the way. This work contains as many of their ideas as mine. I’m particularly grateful to Casey and Julia, who each went well beyond the call of duty in helping me throughout this process. Profound thanks also to the mind and perception group, especially James Gulledge and Emily Prychitko, for extensive support in the formulation of these ideas. This work has been further influenced and supported by an endless number of people, especially Austin Andrews, Renee Bolinger, Gabriel Bruguier, Judith Carlisle, Albert Casullo, Rebecca Copenhaver, Carl Craver, Justin D’Ambrosio, Aaron Elliott, Will Fleischer, Thor Grünbaum, Marta Halina, David Henderson, Colin Klein, Joe McCaffrey, Ross Pain, Anya Plutynski, Susanna Siegel, Brian Weatherson, and Preston Werner. I’m also grateful to the philosophy departments at the University of Nebraska-Lincoln, Washington University in St. Louis, and The Australian National University, for providing the supportive environments that allowed this work to flourish.

Benjamin Henke

Washington University in St. Louis

August 2021

ABSTRACT OF THE DISSERTATION

The Disunity of Perception

by

Benjamin Henke

Doctor of Philosophy in Philosophy-Neuroscience-Psychology

Washington University in St. Louis, 2021

Professor Casey O'Callaghan, Chair

This dissertation argues for disunity in perceptual processing: rather than outputting to a single 'centralized' cognitive system, separate perceptual processing pathways produce different person-level representations for different purposes. I argue that this disunity has important implications for abstract theorizing in the philosophy and cognitive science of perception, for experimental methodology, and for our understanding of the normative role of perception itself. The first three chapters explore models in cognitive neuroscience which support perceptual disunity. In Chapter 1, I argue that a version of Milner and Goodale's (2006) influential 'two visual streams hypothesis' survives its recent empirical challenges. In Chapter 2, I present a new conception of visual streams as core mechanisms for person-level representation-types. In Chapter 3, I argue that the function of functional division in the visual system is to realize both 'task-coupled' and 'task-decoupled' processing approaches in response to the same downstream task. Chapters 4 and 5 begin the work of investigating the normative implications of disunified perception. Chapter 4 argues for a responsibilist account of the epistemology of perception, on which we are epistemically responsible to indirectly influence our perceptual states in order to produce good epistemic outcomes. Chapter 5 explores the role of non-epistemic factors in

determining our epistemic responsibilities. I argue for an account of ‘pragmatic encroachment’ that I call ‘degree encroachment’, on which non-epistemic factors influence the *degree* of continuous justification conferred by one’s perceptual experiences.

Introduction

In "The Modularity of Mind," Jerry Fodor (1983) begins by considering the implications of conceiving of the mind as a computer, that is, as a symbol manipulator that operates only on the syntactic relations between those symbols. Thus, whereas abstract theorizing about computers emphasizes the nature of those syntactic operations, a crucial part of theorizing about the mind is to understand how worldy information is transformed into a format which is accessible to thought—that is, how perceptual and other 'input' systems turn sensory stimulation into syntactically-structured representations for cognition—and how thought itself is transformed into a format which is accessible to action systems—that is, how syntactically-structured representations in cognition are transformed back into analog behaviors. In setting out his modularist theory of mind, he aims to begin to answer that first question: input systems are modules whose primary function is to transform 'analog' perceptual inputs into syntactically structured inputs for thought.

Fodor's presentation begins by considering a (plausible) thesis about what *cognition* is — namely that cognition is a computer that functions to transform perception inputs into action outputs — it then asks what job is performed by *perception*, given this characterization of cognition — namely, it must transform sensory stimulation into a format accessible to thought — and it then concludes that what perception *is* is a system which performs this role. On Fodor's telling, "what perception must do is to so represent the world as to make it accessible to thought…" thus "perception is a mechanism of belief fixation par excellence: the normal consequence of a perceptual transaction is the acquisition of a perceptual belief" (Fodor 1983, 40). This way of approaching the study of perception — by first considering its role in realizing the function of cognition — implies a 'centralized' architecture of the mind on which cognition always mediates perceptual inputs and action outputs, what Susan Hurley (1998, 1) has called "the classical sandwich." And that centralized architecture in turn entails what I'll call a 'unified' conception of perception: which perceptual processes uniformly output to a single downstream node (i.e. cognition).

This dissertation is motivated by the dominance of the foregoing approach to the study of perception and to our corresponding commitment to a unified conception of perception. The unity assumption oversimplifies the nature of perceptual systems and thus their downstream impacts on action. This oversimplification leads us to overlook important aspects of the nature of perception, of our downstream actions, and thus of the fundamental nature of the mind. My aim is therefore to highlight ways in which the mind is *not* unified around a central processor, i.e. that perception is not unified, but instead generates different perceptual products for different downstream tasks. Rather than thinking of the mind as a computer, then, I propose that it's more accurate and more productive to think of the mind as an organized network, with separate but interconnected systems operating in parallel to realize mental functions. Recognizing this, I think, has broad-reaching implications for our conception of ourselves and of the sorts of ways in which we can and should exhibit self-control.

The first aim of this dissertation is to argue for what I call 'representational disunity' in perception, the claim that perceptual systems perform their role by the creation of multiple, distinct representation types for distinct person-level functional roles. My argument for that view is primarily empirical, and turns on facts about how mental systems organize. Establishing representational disunity across perception would require attention to considerations from a broad range of literatures in cognitive science and philosophy. Thus, my goals here are comparatively modest: I explore just a few models from cognitive neuroscience which indicate representational disunity.

My primary aim in the first three chapters of the dissertation is to defend a version of 'the two visual streams' view, most prominently defended by Milner and Goodale (2006), that both survives recent empirical and philosophical challenges and which better orients focus around representational disunity, which I argue is the most important implication of that model.

Chapter 1 defends a new model of the visual streams that I call "Dorsal-Control". According to Dorsal-Control, dorsal stream representations impact action systems directly, without cognitive intermediaries. Since perceptual representations also impact action via cognition, direct dorsal control entails representational disunity. I argue that Dorsal-Control is both consistent with the current empirical evidence and nonetheless has the central implications of the stronger model.

Representational Disunity is neutral between a wide variety of underlying architectures, including ones on which separate representation-types are created by separate visuomotor modules (as Milner and Goodale suggest), and ones on which those representations are generated by more integrated systems (as I suggest in Chapter 1). Given this neutrality, one might think that my conclusions from Chapter 1 leave no place in cognitive science for the notion of a 'stream' of visual processing. My aim in Chapter 2 is thus to re-situate the literature on the visual streams around a conception of the streams as 'core mechanisms' for functionally-individuated representation types. I first argue that substantial inter-stream connection undermines the prospect for a modular conception of the streams. I then articulate a new account of the streams on which they are neurally-isolable core mechanisms for functionally-individuated representation types. I then go on to argue that this conception of a visual stream both survives recent empirical challenges to the visual streams while doing meaningful theoretical and explanatory work in the cognitive science of perception.

The mechanistic conception of a visual stream salvages the visual streams from otherwise apparently devastating objections. In doing so, however, it raises a more fundamental question about the function of the visual streams. Why does the perceptual system divide and conquer, dedicating separate core mechanisms to different roles, rather than simply generating unified representations which are deployed across task-types. This question, addressed in Chapter 3, is at the heart of my rejection of the unified conception of the mind. I begin by presenting evidence against the received view, defended by Milner and Goodale, that the separate visual streams allow different streams to meet the conflicting task-demands of perception and action. Specifically, I argue that there is now evidence for the involvement of multiple streams—with very different computational properties—involved in the same downstream tasks. If so, then the division between the streams cannot be justified by a simple conflict in the computational demands of different tasks. Next, I consider an account, reminiscent of Briscoe and Schwenkler's (2015) account of the role of the two visual stream in motor guidance, according to which the function of the division is instead to realize the *task-neutral* computational advantages of employing both quick-and-dirty and slow-and-intelligent computations. I argue that this view has clear advantages over Milner and Goodale's account, but fails to explain why some visual streams, specifically the dorsal stream and low-road, appear to be largely dedicated to the performance of just one downstream task while the ventral stream appears to operate taskneutrally. I then defend an account according to which the division between streams is in fact one between *task-coupled* and *task-decoupled* computations. This account—specifically the dorsal stream and low-road—are task-coupled, in that they are dedicated to the quick-and-dirty performance of task-specific computations while another system—specifically the ventral stream—is task-decoupled, such that it is engaged in slow and deliberate task-neutral computations. The division between the visual streams has the function of allowing for both task-coupled and task-decoupled computational resources to be applied to the same downstream types of task, optimizing performance.

At the end of Chapter 3, I highlight an important implication of disunity on our capacity for self-control. While we can often realize direct top-down control over ventral stream's input on action, our control over dorsal stream and the low-road is more indirect. This is because the latter systems are *coupled,* preventing synchronous influence over its outputs, but allowing for broader regulation via modulation of ourselves or our environment. In Chapter 4, I begin to explore the normative implications of indirect control. Focusing on perceptual experience, I develop a framework for analyzing the epistemic import of indirect influence on perception. I argue that beliefs based on irresponsibly formed experiences—experiences whose causes were not appropriately regulated by the subject—are doxastically unjustified. This chapter thus sets out a responsibility framework whereby we can analyze the epistemic import of indirect influences on our mental states and downstream actions. This framework allows us to begin to understand our epistemic obligations with respect to coupled perception/action systems.

The argument presented in Chapter 4 relies on the idea that non-epistemic features of one's context can help determine one's epistemic responsibilities. Specifically, it assumes that being an expert perceiver in a domain, in the sense that one is deferred to in making perceptual judgments in that domain, gives you special obligations to follow rules about perceiving in that domain. This account entails a version of pragmatic encroachment—on which non-epistemic factors help to determine epistemic ones. Chapter 5 attempts to articulate a new kind of encroachment that is most fitting with the account in Chapter 4. Specifically, it defends what I call 'degree encroachment', the view that non-epistemic factors can encroach on the degree of continuous justification one has for one's belief. Paired with the responsibilist epistemic account in Chapter 4, degree encroachment suggests that our epistemic obligations concern ways in which we are reasonably expected to comport ourselves in order to realize good epistemic outcomes, even when that influence is indirect. Thus, the normative account developed in Chapters 4 and 5 allow us to begin to unify more traditional normative obligations with those suggested by the disunified mind.

References

Briscoe, R., & Schwenkler, J. (2015). Conscious Vision in Action. *Cognitive Science*, *39*(7), 1435–1467. https://doi.org/10.1111/cogs.12226

Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. MIT press.

Hurley*, S* L. (1998) Vehicles, Contents, Conceptual Structure, and Externalism*. Analysis,* ***58*** *(1)* pp. 1-6

Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action (2nd ed).* Oxford University Press.

Chapter 1: A Fresh Look at the Two Visual Streams

According to what I'll call the "Two Visual Systems Account" (TWO-SYSTEMS), the visual system is divided into two independent sub-systems, a ventral system implementing "vision for perception" and a dorsal system implementing "vision for action" (Milner and Goodale, 2006). TWO-SYSTEMS is widely discussed in philosophy due to the counterintuitive role that it posits for conscious experience in the control of actions. However, recent evidence undermines the model's core tenets: it no longer appears that the ventral and dorsal streams constitute isolated processing systems, and there is now evidence for the involvement of both streams in both conscious experience and online motor control. I articulate a new "Direct Dorsal Control Account" (DORSAL-CONTROL), show that it is immune to three empirical challenges facing TWO-SYSTEMS, and show that it nonetheless has similarly significant implications for the perceiving mind.

Section 1.1 articulates TWO-SYSTEMS and its philosophical implications. Section 1.2 describes three empirical challenges to TWO-SYSTEMS. Section 1.3 articulates DORSAL-CONTROL and shows that it is immune to the three challenges. Section 1.4 articulates DORSAL-CONTROL's philosophical implications.

1.1 The Two Visual Systems Account

The primate cortical visual system is composed of two anatomically distinct streams, a ventral stream extending laterally from V1 into inferotemporal cortex, and a dorsal stream extending dorsally from V1 into the posterior parietal lobe. TWO-SYSTEMS holds that these streams constitute independent sub-systems, each with a proprietary function.

To explain this proposal, Goodale and Milner (2004) draw an analogy to the guidance of a hypothetical Martian rover. The rover's job is to navigate the terrain and collect samples. To help it accomplish these goals, engineers on earth assess the contents of the rover's sensors and send action programs about what it should do next. But since there is a substantial delay between Earth and Mars, it would be inadvisable for the rover to execute exact movements on the basis of mission control's program. Instead, the rover must update the program in response to in-the-moment conditions. Thus, the rover is fitted with an additional onboard system which independently assesses the scene and makes fine-grained adjustments to the action program. When collecting samples, for example it uses its sensors and onboard systems to guide its limbs toward the target. The rover's behavior is thus controlled by two systems: a smart but slow controller system which *selects* action programs, and a dumb but quick onboard system which *guides* the fine-grained execution of those programs.

TWO-SYSTEMS proposes a similar division in the cortical visual system. The ventral system outputs directly to cognition, enabling smart but slow selection of action programs. And the dorsal system outputs directly to action systems to enable dumb but quick guidance of those programs. For example, if I aim to catch a baseball hit into the outfield, ventral stream and downstream cognition select an action program—e.g. "run over there and catch the ball"—while the dorsal stream guides the execution of that program—e.g. by enabling me to keep a constant angle between the horizon and the ball, ensuring that I arrive at the bottom of the ball's arc as it moves across the field.

TWO-SYSTEMS posits ventral and dorsal *systems* which perform their respective functions in isolation. It follows both that cognition employs exclusively ventral stream representations in the selection of action programs and that action systems are guided exclusively by dorsal stream representations in the execution of those programs. Three important implications follow from this conception of the two visual streams.

First, TWO-SYSTEMS controverts pretheoretic assumptions about conscious experience's role in action. According to what Clark (2001) calls the "assumption of experience-based control," conscious visual experience is employed in "the control and guidance of fine-tuned real world activity". But TWO-SYSTEMS holds that online motor control is instead executed using unconscious dorsal stream processes, not conscious ventral stream ones.

Second, TWO-SYSTEMS has revisionary architectural implications. A popular view of the mind holds that perceptual processes function to transform sensory information into a format accessible to central cognitive processes (Fodor, 1983; for an opposing view, which is much closer to the one defended here, see Nanay, 2013). But if dorsal stream processes influence action systems *directly*, without cognitive intermediaries, then this common assumption is false. Instead, perceptual systems perform multiple distinct functions in perceptual/motor control, only some of which directly involve cognition.

Third, TWO-SYSTEMS informs the role of action in the study of conscious experience. Researchers often infer that a feature is present in perceptual experience from the fact that our behavior is sensitive to that feature. But if action systems employ unconscious dorsal stream processes in the online control of behavior, then such inferences are suspect. For example, Clark (2001) has argued that TWO-SYSTEMS undermines common arguments for so-called "cognitive overflow," the claim that the contents of perceptual experience are more detailed than our cognitive representations of them.

1.2 Three Empirical Challenges to TWO-SYSTEMS

TWO-SYSTEMS claims that ventral and dorsal streams constitute independent processing systems and that each has a separate, proprietary function. Each of these claims is challenged in the empirical literature.

First, empirical evidence undermines TWO-SYSTEMS's claim that the ventral and dorsal constitute independent processing systems. While exact conditions for independence are controversial, it is widely agreed that significant functional or informational connectivity between streams is evidence against independence. And there is now evidence for widespread interconnectivity of both types between the two streams. First, neuroanatomical studies have revealed substantial functional connectivity between ventral and dorsal streams (for reviews, see van Polanen and Davare, 2015; Cloutman, 2013; Grafton, 2010). And this functional connectivity appears to produce informational connectivity as well, such that information processed in one stream is employed in computations in the other stream (Schenk and McIntosh, 2010; van Polanen and Davare, 2015). In sum, as Schenk and McIntosh (2010) put it "The ubiquity and extent of inter-stream interactions suggest that we should reject the idea that the ventral and dorsal streams are functionally independent processing pathways."

Second, there is now evidence for dorsal stream influence on conscious experience. As Schenk and McIntosh (2010) report, while early studies of optic ataxics (who have isolated dorsal stream damage) appear to demonstrate deficits in perceptual-motor control without corresponding deficits in conscious report, apparently confirming TWO-SYSTEMS's predictions (Karnath & Perenin, 2005), more refined studies have found that optic ataxics have subtle but persistent differences in conscious experience (Michel and Henaff, 2004; Perenin and Vighetto, 1988; Pisella et al., 2009; Striemer *et al.*, 2007, Striemer, Chapman, & Goodale, 2009; Striemer *et al.*, 2009). This evidence suggests that dorsal stream processes influence either conscious object representations or conscious attention or both (van Polanen and Davare 2015). Philosophers have also suggested that dorsal stream processes may influence structural features of conscious experience, such as its spatial format (Wu, 2014) or 'feeling of presence' (Matthen, 2005).

Third, there is also evidence for ventral stream's involvement in online motor control. Here again, while early lesion studies among those with isolated ventral stream damage seemed to support TWO-SYSTEMS's predictions—demonstrating deficits in conscious report without corresponding deficits in perceptual-motor control—more refined studies have revealed that such subjects have systematic deficits in online motor control when performing complex tasks, such as fitting a T-shaped object through a slot (Goodale et al., 1994; McIntosh and Lashley, 2008; Schenk and McIntosh, 2010). Briscoe and Schwenkler (2015) have suggested a kind of dual-systems account of the role of ventral and dorsal streams in perceptual motor control: dorsal stream is responsible for the direct control of easy, familiar aspects of tasks, but heavily recruits ventral stream representations in executing novel or difficult motor commands.

These three challenges, taken together, constitute a complete refutation of TWO-SYSTEMS. The rest of this chapter argues that a new view, which captures the central functional claims of Milner and Goodale's account, and which has similarly important implications for the perceiving mind, survives the challenges.

1.3 The Direct Dorsal Control Account

The central claim of this chapter is this: the main upshots of Milner and Goodale's model hold if there are perceptual representations which directly influence action systems, without cognitive intermediaries. DORSAL-CONTROL is the claim that at least some dorsal stream representations play this role. In this section, I clarify DORSAL-CONTROL and show that it is immune to the empirical challenges raised against TWO-SYSTEMS.

The claim that a representation influences action systems *directly,* without cognitive intermediaries, is equivalent to the claim that action systems have direct access to those representations, not just downstream cognitive representations. Thus, DORSAL-CONTROL holds that action systems directly recruit dorsal stream representations in the execution of action programs. This claim follows from TWO-SYSTEMS, since the latter entails that the dorsal system implements "vision for action" without cognitive intermediaries. But DORSAL-CONTROL is strictly weaker than TWO-SYSTEMS, as it entails neither that the ventral and dorsal streams constitute independent processing systems nor that either has a proprietary functional role.

It's important to distinguish DORSAL-CONTROL from the claim that dorsal stream representations are *unique* in their guidance of action programs. The former holds only that dorsal representations at least sometimes perform a particular functional role, not the stronger claim that they are the only perceptual representations to do so. It is nonetheless worth noting that the stronger claim is also plausibly true. That is, even if, as the third challenge contends, ventral stream processes influence action programs, it is likely that they do so only via downstream cognitive representations and processes. Thus, DORSAL-CONTROL provides a framework via which one can rescue the central functional division proposed by Milner and Goodale.

DORSAL-CONTROL is a claim about the *downstream effects* of dorsal stream representations, not its upstream causes. Thus, it is consistent with the possibility of causal influence of ventral stream or cognitive representations and processes on dorsal stream processes. Even if, for example, dorsal stream processes recruit ventral stream object representations in carrying out complex guidance tasks (van Polanen and Davare, 2015), the resulting dorsal stream representations can still influence action systems directly without cognitive (or ventral stream) intermediaries.

Moreover, DORSAL-CONTROL makes no claim of exclusivity in dorsal stream's effects. That is, the claim that some dorsal stream representations directly influence action programs is consistent with those or other dorsal stream representations also influencing states and processes elsewhere in the mind. Thus, DORSAL-CONTROL is consistent with dorsal stream influences, even conscious influences, on ventral stream or cognition.

Consider the first challenge. Since TWO-SYSTEMS posits independent ventral and dorsal processing systems, it is challenged by evidence for substantial functional and informational connectivity between the streams. But since DORSAL-CONTROL doesn't require that the streams be independent, it is consistent with such connectivity. What matters is that dorsal stream representations directly influence action programs. And this claim is widely accepted in the literature, even among TWO-SYSTEMS's opponents. As Schenk and McIntosh (2010, 52) put it in one of their many concessive remarks on this point: "[Milner and Goodale's account] captures some broad patterns of functional localization, but… the specializations of the two streams are relative, not absolute."

Next, consider the second challenge. Since TWO-SYSTEMS holds that ventral stream is exclusively responsible for the formation of conscious experience, it is challenged by evidence for dorsal stream influence on conscious experience. DORSAL-CONTROL, by contrast, makes no claim about ventral stream's exclusive role. It is thus consistent with dorsal stream influence on conscious experience.

One might reply that if dorsal stream contributes to conscious experience, then on the reasonable assumption that a state is conscious only if it is accessible to cognition, dorsal stream must output to cognition, not directly to action programs. This reply is mistaken in two ways. First, it assumes that dorsal stream can influence conscious experience only if some of its representations are themselves conscious. But dorsal stream can influence conscious experience —even if its representations are unconscious—by having conscious representations among its *downstream effects*. This would happen, for example, if dorsal stream had downstream causal influence on conscious ventral stream representations. This possibility is inconsistent with TWO-SYSTEMS, since that theory holds that ventral stream is causally isolated from dorsal stream and that ventral stream has *proprietary* influence over conscious visual representations. But it is consistent with DORSAL-CONTROL, which makes neither claim.

Second, the reply assumes that a representation's being directly accessible to cognition is inconsistent with that representation also being directly accessible to action systems. But a single representation can be recruited by multiple distinct processes or systems. Thus, it's consistent with DORSAL-CONTROL—but not, again, with TWO-SYSTEMS—that the very representations which directly guide action programs also directly contribute to conscious experience. As we'll see below, the latter does not negate the importance of the former.

Finally, consider the third challenge. Since TWO-SYSTEMS holds that dorsal stream is exclusively responsible for online motor guidance, it is challenged by evidence for ventral stream influence on online motor activity. But DORSAL-CONTROL simply claims that dorsal stream representations are directly involved in online motor guidance, not that that involvement is exclusive. It is thus consistent with ventral stream involvement. Indeed, those who posit ventral stream influence on online motor activity either posit such influence *via* dorsal stream representations or, as Briscoe and Schwenkler (2015) suggest, in conjunction with it. Thus, ventral stream influence on online motor activity is not a threat to DORSAL-CONTROL.

1.4 Implications of DORSAL-CONTROL

To see why DORSAL-CONTROL entails the same general implications for the perceiving mind, it's helpful to see that the three implications discussed in Section 1 are intricately linked. Our pretheoretic assumption is that experience plays a direct role in the fine-grained guidance of our actions. And we assume that this role is mediated by our choices. That is, we assume that we use fine-grained details of our perceptual experiences to make decisions about how we will enact our will in the world. But if TWO-SYSTEMS is right, then both assumptions are false: if unconscious dorsal stream processes uniquely guide online motor activity, without cognitive intermediaries, then experience and downstream cognition have a less direct role in that activity. And these first two implications entail the third: if experience and our decisions play only an indirect role in online control, then direct inferences from behavior to experience (or cognition) are suspect.

Because DORSAL-CONTROL, like TWO-SYSTEMS, breaks the intricate connection previously assumed between conscious experience and motor control, a connection which we assumed was mediated by cognition, it has these same implications. That is, if dorsal stream representations influence online motor control directly, then cognition-mediated perceptual representations do not do the work we thought they did in controlling our actions. And this undermines direct inferences from behavior to experience.

One might respond as follows: "You acknowledged above that it's consistent with DORSAL-CONTROL that the very dorsal stream representations which directly influence online motor activity are also consciously accessible. Don't the implications rely on a lack of conscious accessibility? That is, isn't it an important feature of TWO-SYSTEMS that *unconscious* dorsal stream representations guide online motor activity?" This, I think, is the central mistake we make about the philosophical import of the two visual streams. It doesn't so much matter whether dorsal stream representations are conscious. It matters whether they control behavior *because* they're conscious. And according to DORSAL-CONTROL, they do not: even if dorsal stream representations are consciously accessible, they perform their guiding role *directly*, independently of any conscious influence on cognition.

An analogy will help make this final point. Imagine driving a car which has screens and cameras in place of windows. As you move along, you look at the screens and move the wheel and pedals as you normally would. It *seems* to you that you have the same kind of control over the car as you normally do. That is, you use what you see on the screen to make decisions about how you move the wheel and pedals to get you where you want to go. But this impression is mistaken. The car is instead driven by an autopilot. This autopilot uses your movements to roughly determine where it will go (this is analogous to the role of ventral stream and cognition in *selecting* action programs), but the actual online control of the car is entirely controlled by its independent decisions (analogous to the direct role of dorsal stream in the *guidance* of action programs).

The picture given to us by TWO-SYSTEMS is one on which the system creating the images on the screen is entirely independent of the system which controls the car. This is analogous to the idea that the dorsal stream representations controlling online activity are unconscious. But suppose instead that the autopilot operates over the very representations you employ in making your decisions. That is, the autopilot, like you, is looking at the images on the screen and making decisions about how to maneuver the car in light of those images. You're looking at *the very representations* employed in the control of the car. This picture is analogous to one on which dorsal stream representations both directly influence online motor control (and thus DORSAL-CONTROL is true) and are consciously accessible to cognition. Does this dual access salvage our pretheoretic judgments about the perceiving mind?

I think not. The representations that we see on the screen are employed in the online control of the car, but *not by us*. The representations on the screen control the behavior of the car *and* we can see them, but they do not control the behavior of the car *because* we can see them. That is, even on this picture, the autopilot is performing a role we don't expect it to perform. This picture is just as surprising as, and more empirically plausible than, the one suggested by TWO-SYSTEMS.

References

Briscoe, R., & Schwenkler, J. (2015). Conscious Vision in Action. *Cognitive Science*, *39*(7), 1435–1467.

Clark, Andy. 2001. Visual Experience and Motor Action: Are the Bonds Too Tight? *The Philosophical Review*, *110*(4), 495–519.

Cloutman, L.L.. (2013). Interaction between dorsal and ventral processing streams: where, when and how? Brain Lang. 127, 251–263.

Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology.* MIT press.

Goodale, M. and Milner, A. (2004). *Sight Unseen: An Exploration of Conscious and Unconscious Vision.* Oxford University Press.

Goodale, M. A., Jakobson, L. S., Milner, A. D., Perrett, D. I., Benson, P. J., & Hietanen, J. K. (1994). The nature and limits of orientation and pattern processing supporting visuomotor control in a visual form agnosic. *Journal of Cognitive Neuroscience*, *6*, 46–56.

Grafton, S.T., (2010). The cognitive neuroscience of prehension: recent developments. *Exp. Brain Res.* 204, 475–491.

Karnath, H. O., & Perenin, M. T. (2005). Cortical control of visually guided reaching: Evidence from patients with optic ataxia. *Cerebral Cortex*, *15*, 1561–1569.

Matthen, M. (2005). *Seeing, doing, and knowing: A philosophical theory of sense perception*. Oxford University Press.

McIntosh, R. D., & Lashley, G. (2008). Matching boxes: Familiar size influences action programming. Neuropsy- chologia, 46, 2441–2444.

Michel, F., & Henaff, M. A. (2004). Seeing without the occipito-parietal cortex: Simultagnosia as a shrinkage of the attentional visual field. Behavioural Neurology, 15, 3–13.

Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action* (2nd ed). Oxford University Press.

Nanay, B. (2013). *Between perception and action*. Oxford University Press.

Perenin, M. T., & Vighetto, A. (1988). Optic ataxia: A specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain*, *111*, 643–674.

Pisella, L., Sergio, L., Blangero, A., Torchin, H., Vighetto, A., & Rossetti, Y. (2009). Optic ataxia and the function of the dorsal stream: contributions to perception and action. *Neuropsychologia*, *47*(14), 3033-3044.

Schenk, T., & McIntosh, R. D. (2010). Do we have independent visual streams for perception and action? *Cognitive Neuroscience*, *1*(1), 52–62.

Striemer, C., Blangero, A., Rossetti, Y., Boisson, D., Rode, G., Vighetto, A., et al. (2007). Deficits in peripheral visual attention in patients with optic ataxia. *Neuroreport*, *18*, 1171–1175.

Striemer, C. L., Chapman, C. S., & Goodale, M. A. (2009). "Real-time" obstacle avoidance in the absence of primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 15996–16001.

Striemer, C., Locklin, J., Blangero, A., Rossetti, Y., Pisella, L., & Danckert, J. (2009). Attention for action? Examining the link between attention and visuomotor control deficits in a patient with optic ataxia. *Neuropsychologia*, *47*, 1491–1499.

van Polanen, V., & Davare, M. (2015). Interactions between dorsal and ventral streams for controlling skilled grasp. *Neuropsychologia*, *79*, 186–191.

Wu, W. (2014). Against Division: Consciousness, Information and the Visual Streams. *Mind & Language*, *29*(4), 383–406.

Chapter 2: Visual Streams as Core Mechanisms

In Chapter 1, I defended DORSAL-CONTROL, the claim that dorsal stream representations directly influence action systems, without cognitive intermediaries. This thesis, along with the claim that visual experiences influence action via cognition, entails *representational disunity,* the claim that multiple distinct representations are involved in performing vision's person-level functional roles.

In defending DORSAL-CONTROL, I argued that it, but not a dual-systems view, is consistent with evidence for substantial functional interaction between the streams, and for each stream's involvement in the other's characteristic function. But I also argued that DORSAL-CONTROL retains the central implications of Milner and Goodale's (2006) model. Such an argument might appear to leave little room for the idea of a visual stream in the cognitive science of perception. If what matters is simply a disunity in representations—which could be realized in an otherwise unified processing system—then it's no longer clear what work is left for the idea of separate processing streams for different purposes.

The aim of this chapter is to re-situate the idea of a visual stream in the study of perception. Visual streams play an explanatory role, I'll argue, but not the one traditionally ascribed to them among defenders of the two visual streams. I suggest two revisions to that traditional notion. First, I'll argue that we should conceive of a visual stream, not as an isolated processing module, but rather as the core *mechanism* for a functionally-individuated representation-type. Second, and in virtue of this mechanistic conception of a visual stream, I'll suggest that the study of visual streams is *explanatorily secondary* to the study of representational disunity. That is, the central explananda in the study of the visual streams are neither anatomical and computational features of visual processing nor downstream behavior. Rather, the central explananda are the existence and features of the representations the visual streams characteristically produce. These revisions are both conservative—in that they capture the central work required of a visual stream in the empirical literature—and revisionary—in that they reject the conceptual apparatus typically employed in that literature.

The chapter will proceed in three sections. In Section 2.1, I'll articulate the traditional conception of the streams as modules, and provide two reasons—one empirical and one theoretical—for why we should reject that conception. In Section 2.2, I'll present a new conception—on which a visual stream is a core mechanism for a functionally-individuated representation type—and highlight its consistency with the empirical evidence. In Section 2.3, I'll articulate the role that this new conception of a visual stream can play in the study of visual perception.

2.1 The Visual Streams as Modules

In *The Modularity of Mind,* Jerry Fodor (1983) lists nine features that characterize modules. Modules are *informationally encapsulated* — in that they operate over only local information stores — they're *inaccessible* — in that other systems do not have access to those information stores — they're *domain specific* — in that they operate over narrow subject matters — their processes are *mandatory, fast*, and have '*shallow*' (roughly, coarse-grained) outputs, and they're underwritten by a *fixed neural architecture* with *characteristic and specific breakdown patterns* and a *characteristic ontogenetic pace and sequencing*. Fodor holds that a system is a module if it has these properties "to some interesting extent" (Fodor, 1983, 37). Moreover, among the nine, Fodor thought that informational encapsulation was most important, as substantial informational encapsulation is both essential for modularity and explanatorily prior to many of the other features. While there has been substantial debate over both his account of modularity and the empirical case for mental modules, his treatment, and in particular his focus on informational encapsulation, has remained the default working account of modularity in philosophy and the cognitive sciences.

At the beginning of their seminal book on the two visual streams, Milner and Goodale (2006) contrast prior 'input-only' conceptions of modularity in the visual system — such as Marr's (1982) postulation of visual primitives — with their own account, which also posits modularity 'on the output side' (Milner and Goodale 2006, 6, 14). That is, where prior accounts posited visual modules which function to process local features of an ultimately unified visual representation, Milner and Goodale argue that the visual streams constitute 'visuomotor' modules, which generate separate representations for different downstream tasks, i.e. visual perception and motor control (2006, 5). In beginning their book in this way, they thus cement their commitment both to what I've called representational disunity, and to a modularist account of that disunity.

Milner and Goodale note that *sub-cortical* visuomotor modules are present throughout the animal kingdom. Ingle (1973), for example, traced separate subcortical pathways in frogs (*Rana pipiens*) from the retina to the motor nuclei for different kinds of visually guided behavior. "The fact that the frog possesses this parallel set of independent visuomotor pathways," Milner and Goodale say, "does not fit well with the common view of a visual system dedicated to the construction of a unified representation of the external world" (2006, 11). Moreover, they point to evidence for a similar division in the subcortical visual system of rats, concluding "the modular organization of visuomotor behaviour in representative species of at least one mammalian order, the rodents, appears to resemble that of much simpler vertebrates such as the frog and toad. In both groups of animals, visually elicited orienting movements, visually elicited escape, and visually guided locomotion around barriers are mediated by quite separate pathways from the retina right through to motor nuclei in the brain-stem and spinal cord. This striking homology in neural architecture suggests that modularity in visuomotor control is an ancient (and presumably efficient) characteristic of vertebrate brains" (2006, 17).

The questions that Milner and Goodale seek to answer, then, are, first, does visuomotor modularity extend into the cortical visual system? and, second, is there evidence for such modularity in 'higher' mammalian orders, such as primates, whose cortical visual systems are more developed? They answer 'yes' to both questions. Specifically, they claim that "two separate networks of areas have evolved in the primate visual cortex: a perceptual system which is indirectly linked to action via cognitive processes, and a visuomotor system which is intimately linked with motor control" (2006, 20).

Given the dominance of Fodor's conception of a module, however, the preceding discussion should puzzle us. In explicitly arguing for visuomotor modularity, Milner and Goodale focus entirely on evidence for the *anatomical* separation between visual subsystems and the functions associated with those subsystems. But to establish that the streams are modules, one must show that they satisfy certain *computational features* — especially that they operate only over local information stores — not simply that they're anatomically separable.

Milner and Goodale clearly think of the visual streams as broadly Fodorian modules, however, and they provide evidence which supports a number of Fodor's criteria, beyond mere anatomical separation. The initial evidence for their account, for example, comes from lesion studies in which subjects with isolated damage to one stream impacts that stream's characteristic function (visual perception or motor control) without impacting the other stream's function (Milner & Goodale, 2006). This suggests characteristic and specific breakdown patterns of the two streams. And Milner and Goodale clearly take these breakdowns as evidence for informational encapsulation (e.g., 2006, 121). Similarly, their account is famously supported by studies demonstrating that visual perception, but not motor control, is subject to certain illusions, such as the Ebbinghaus or Müller-Lyer illusions (2006, ch. 6). This suggests that the two streams operate over separate information stores, supporting informational encapsulation and inaccessibility. Finally, Milner and Goodale argue extensively that each stream is specialized to certain contents — such as joint angles for motor guidance or object identities for visual perception — and that they operate over different spatial formats — egocentric formats for motor guidance and allocentric ones for perception (e.g. 2006, section 4.1). This suggests a kind of domain-specificity between the two streams.

Milner and Goodale's claim that the two streams are dedicated to particular functions — visual perception and motor guidance, respectively — has been largely accepted in the empirical literature. The debate regarding their account has focused, instead, on whether the streams constitute informationally-isolated modules. As Schenk and McIntosh (2010, 52) put it in their review of evidence against the model, "…the *perception–action* model captures some broad patterns of functional localization, but the specializations of the two streams are relative, not absolute. The ubiquity and extent of inter-stream interactions suggest that we should reject the idea that the ventral and dorsal streams are functionally independent processing pathways."

Thus, both supporters and detractors of Milner and Goodale's model seem to agree that the visual pathways are interesting to the extent that they constitute isolated modules. By extension, it's also generally agreed that the model and its implications rise or fall with the claim that the streams are at least largely informationally encapsulated.

This conception of the visual streams faces two challenges. The first, discussed in Chapter 1, is empirical: there is now evidence for substantial informational connectivity between the streams and with extra-perceptual systems. The dorsal stream appears to recruit ventral stream representations to enable sensitivity to allocentric spatial features (Schenk & McIntosh, 2010), pictorial depth (Schenk & McIntosh, 2010), and object identity (Carey, Harvey, & Milner, 1996; Dijkerman, Milner, & Carey, 1998; McIntosh, Dijkerman, Mon-Williams, & Milner, 2004; Dijkerman, McIntosh, Schindler, Nijboer, & Milner, 2009; van Polanen & Davare, 2015) when engaging in tasks requiring complex fine-tuning of grasp. And the ventral stream appears to recruit dorsal stream representations to determine object location and orientation (Michel & Henaff, 2004; Perenin & Vighetto, 1988; Pisella et al., 2009) and to control conscious attention (Michel & Henaff, 2004; Pisella et al., 2009; Striemer *et al.,* 2007; Striemer, Chapman, & Goodale, 2009; Striemer *et al.*, 2009). Such connectivity between the streams directly challenges the claim that the streams are informationally encapsulated and inaccessible modules.

The second problem with the modularist conception of the visual streams is more theoretical. As I've suggested above, the empirical debate surrounding the two visual streams has largely centered on the question of whether the visual streams operate over proprietary information stores and thus are modular. This trend gives the impression that the central implications of the model rise or fall with modularity. But most of Milner and Goodale's most radical claims are independent of the modularity question. I'll review several such instances.

First, of course, the central claim of Milner and Goodale's model is a claim about the central function of the respective streams. The dorsal stream has the function of guiding online motor activity while the ventral stream has the function of generating largely conscious perceptual representations. That claim can be true even if the streams are not modules. This functional division might reflect the claim, defended in Chapter 1, that the streams are vehicles for the representations driving these respective tasks. This would constitute a minimal functional division between the streams. Moreover, even if the streams interact, it could still be the case that the bulk of the processing determining the different features of these different representations are also performed within the respective streams. For example, van Polanen and Davare (2015) review evidence that dorsal stream processes recruit ventral stream object representations when engaging in tasks requiring complex fine-tuning of grasp, for example when determining the appropriate grip for tool use. One might therefore appropriately conclude that the ventral stream is involved in motor control in such cases. But van Polanen and Davare themselves suggest that object information is probably not employed *directly* in determining motor control representations, but is rather translated into a format more appropriate for that task. That translation is likely performed within the dorsal stream (van Polanen & Davare 2015, 188; Binkofski & Buxbaum, 2013). Thus, one can coherently hold that the ventral stream influences motor control representations in such cases while also holding that the computations which determine the central feature of those representations are performed within the dorsal stream.

This would remain a sense in which the respective streams are dedicated to their respective functions, even as they interact. Such a claim is falsifiable, of course, but the point is that it is not undermined by the existence of non-proprietary information stores.

Second, as argued in Chapter 1, Milner and Goodale's hypothesis has three crucial theoretical implications: it controverts pre-theoretic assumptions about the role of consciousness in action control, it suggests a revision to 'centralized' conceptions of mental architecture, and it undermines the role of behavior in the study of conscious experience. In Chapter 1, I argued that these implications follow from DORSAL-CONTROL, regardless of whether the perceptual streams constitute isolated perceptual modules. Thus, an emphasis on modularity misleadingly suggests that these implications require it.

Finally, the two streams hypothesis is widely discussed in philosophy due to its apparent commitment to the 'Zombie Action Hypothesis', the claim that the representations driving online motor control are unconscious. The zombie action hypothesis seems inconsistent with our pre-theoretic judgements, which seem to hold that we use our conscious experiences in controlling our actions. Thus, the philosophical discussion surrounding the two streams has centered on whether the evidence entails the zombie action hypothesis. And the question of whether the evidence supports zombie action has widely been taken to be the same question as whether the streams are independent modules for different functions (e.g. Mole, 2009). The trouble with this line of thinking is that the question of modularity and the question of zombie action are not the same. Each can be true without the other: One might have two modular systems which each output conscious representations to cognition. And there could be a single system which produces both conscious representations for perception and unconscious representations for online motor control.

In summary, then, a modular conception of the two streams should be rejected on the grounds that (1) mounting empirical evidence weighs against the streams being modules, and (2) treating the two streams hypothesis as rising or falling with modularity conflates several different issues raised by Milner and Goodale's work. Thus, we should aim for a reorientation of our theoretical conception of the two streams which better accords with the empirical evidence and which better situates the various threads of Milner and Goodale's theory in relationship to one another. I turn now to developing an alternative account which can do this work.

2.2 The Visual Streams as Mechanisms

My aim in the next two sections is to articulate a new conception of the visual streams. That conception must accomplish two goals. First, it must be broadly consistent with the extant empirical evidence. In particular, it must be consistent, unlike the modularity account it replaces, with the now substantial evidence for interaction between the streams and with goings-on elsewhere.

Second, on a new conception, visual streams must do meaningful explanatory and theoretical work in cognitive science that broadly accords with current scientific usage. Of course, a new account cannot completely accord with current usage, which is predominantly modularist. But a new account should help us understand why scientists posit visual streams where they do, and what utility positing such streams has in scientific theorizing. In particular, an account of the visual streams should capture both the role that streams play in explaining empirical results and their broader role in psychological theorizing.

I'll take up each of these goals in turn. In this section, I'll articulate my account and show that is consistent with the extant evidence. In the next section, I'll articulate the role, on my account, that visual streams play in scientific theorizing, showing how it accords with the bulk of extant scientific practice.

Without further ado, my account is that a visual stream is the neurally-individuable core mechanism for the generation of representations with a characteristic person-level functional role.

The central difference between my account and the dominant one is that I claim that a visual stream is merely a *mechanism*, while the dominant one holds that a stream must be a *module*. The former is strictly weaker than the latter. The notion of a mechanism is of a set of entities and their activities which are organized so as to bring about a phenomenon of interest (Craver, 2007). The mechanism which allows me to shift gears on my bicycle, for example, is composed of entities — levers, cables, and derailleurs — and their activities — articulating the lever, applying tension to the cables, turning the primary mechanism on the derailleur, etc. — which collectively realize the phenomenon of interest — pulling the lever increases tension on the cables, which operates the derailleur, which pushes the chain up or down the gear set.

How do we determine the underlying mechanism for a given phenomenon? First, it's important to understand that mechanisms do not necessarily reflect 'joints in nature,' but are determined in part by our interests. This interest-relativity comes in at several points in the identification of a mechanism. In claiming that the mechanism for gear shifting is composed of levers, cables, and derailleurs, for example, I begin by making a choice about a phenomenon of interest, whose precise specification may vary. I could ask, for example, how I shift my front gears versus my back gears. Or, I could ask how gear shifting works on bikes in general, rather than my bike. And so forth. Each of these decisions will result in different specifications of the underlying mechanism.

Second, in determining the components of a mechanism, we begin by making implicit or explicit decisions about where the mechanism begins and ends (i.e. what its inputs and outputs are), and determine components against certain understood background conditions. In identifying the mechanism for gear shifting on my bicycle, it's taken for granted that the bicycle has a rider, which moves the levers, initiating the change. We chose not to include features of the rider (such as the operation of their muscles and joints) in articulating the mechanism for gear shifting. We similarly made certain basic assumptions about the background conditions of the bike. The levers, cables, and derailleurs are each attached to a frame, whose stability is required for the mechanism's operation. The frame's stability is thus a background condition for the operation of the mechanism.

But while decisions about inputs, outputs, and background conditions are interest-relative, they're not necessarily arbitrary. The choices described above, for example, might reflect our interest in bicycle repair: a bicycle lacking a rider isn't broken, and one lacking a stable frame has bigger problems than mere gear shifting. Similarly, mechanistic explanation in science is interest relative, but not necessarily arbitrary, as there may be a wide range of local constraints which help to identify a phenomenon of interest, where a mechanism should begin and end, or what its background conditions should be. Relative to these decisions, mechanistic explanation is paradigmatically ‘objective’: the entities and activities which compose the mechanism are those which causally mediate its inputs and outputs (Craver, 2007; Prychitko, 2019).

Importantly, while all modules are mechanisms, not all mechanisms are modules. Mechanisms needn’t be informationally encapsulated, since they may operate over inputs from outside their boundaries at one or more points in their operation. Mechanisms can share their representations with other operations, and thus needn’t be inaccessible. Mechanisms needn’t be mandatory or fast. And mechanisms needn’t be domain specific, both because a mechanism can itself operate over multiple domains, and because a component of a mechanism may itself feature in multiple mechanisms. Thus, the requirements for being a mechanism are strictly weaker than those of being a module.

Nonetheless, mechanistic explanation can substantially illuminate a phenomenon of interest. Most relevant for our purposes, mechanistic explanation can help us understand how features of a mechanism’s components and their interactions (such as the representational format over which they operate) give rise to features of the target phenomenon (such as a target representation’s format). Craver (2007) argues that mechanistic explanation is the predominant form of explanation in contemporary neuroscience.

The mechanism explaining gear shifting is *local*, in the sense that it involves a relatively constrained set of entities at close spatial proximity with a relatively simple set of interactions. But not all mechanisms are local in this way. The mechanism of climate change is widely distributed, involving many overlapping systems at wide spatial distribution with complex sets of interactions. Neural processes are paradigmatically *global*, in that any given mental phenomenon is typically underwritten by many sets of overlapping systems which interact in highly complex ways to bring about a phenomenon. It is, however, an important part of neuroscientific work to break up such processes into relatively local mechanisms which realize core functions. We say, for example, that the hippocampus is the central hub of the mechanism for episodic memory because the central features of episodic memory are determined by features of hippocampal activity (Tulving, 1983). When a relatively local mechanism explains the central features of a phenomenon of interest, I'll call this the *core mechanism*. Thus, hippocampal activity plausibly constitutes the core mechanism for episodic memory. It's important to note, however, that since not all phenomena are such that their central features are explainable by a single local mechanism, not all phenomena have a core mechanism. Plausibly, for example, visual attention is controlled by many sets of overlapping systems with no one local mechanism explaining the bulk of its central features (Knudsen, 2007). If so, then there is no core mechanism for visual attention. Thus, positing a core mechanism for a given phenomenon is falsifiable; and identifying core mechanisms is theoretically interesting.

On the account sketched above, mechanisms are defined relative to a phenomenon of interest. In the case of the visual streams, the phenomenon of interest is the transformation of retinal stimulation into representations with certain characteristic person-level functional roles. In making this specification, I aim to capture the cases to which researchers have applied the term 'visual stream'. Without the specification that the target representations should have *person-level* functional roles, the definition would entail that there are a wide plethora of sub-personal visual streams: a visual stream for edge detection, one for motion detection, etc. In fact, however, researchers speak of streams only when a division in processing produce person-level functional differences. While that choice is interest-relative, I do not, again, think it is arbitrary: we have special interest in streams that produce person-level representational disunity, because, as argued in Chapter 1, person-level representational disunity has important theoretical and methodological implications.

The account sketched so far, that a visual stream is the core mechanism for the generation of representations with a characteristic person-level functional role, provides only individuation conditions for a single stream. Current usage, however, posits visual streams in multiples. The idea of talking about 'the two visual streams' is that there are at least *two* of them, with different functions. Thus, capturing usage requires that an account also provide differentiation conditions for multiple streams. My account claims that a core mechanism is a visual stream only if it is neurally isolable from some other core mechanism which independently meets the other conditions for being a visual stream.

To conclude the outline of my account, then, the reason that the ventral and dorsal pathways constitute distinct visual streams is that (1) they each function to generate representations with a characteristic person-level functional role (action and cognition guidance, respectively), (2) they each house the core mechanism for the generation of those representations, and (3) these core mechanisms are neurally isolable from one another.

Claim 1 is defended in Chapter 1. Claim 3 simply holds that the ventral and dorsal streams constitute different areas of cortex (that may or may not causally interact). That's not up for dispute. My remaining work for this section, then, is to defend claim 2.

Claim 2 says that the central features of the target representations are determined by goings-on within the respective streams. This is not to say that the streams are independent or that goings-on elsewhere fail to constitute important inputs to the core mechanism. It's just that understanding the goings-on in each stream is the central task required to understand the target representation.

Fortunately, much of the evidence discussed in the two streams literature directly supports this claim. Since that evidence is quite far-ranging, I'll divide the evidence into three types — behavioral studies in neurotypical subjects, lesions studies, and computational analyses — and examine the role of that evidence in supporting claim 2.

Let's begin with behavioral studies in neurotypical subjects. A wide range of studies suggest that a number of visually-guided behaviors, such as reaching, pointing, or grasping, are less sensitive to certain visual illusions than is visual perception (Goodale and Milner, 2018; though see Christiansen, Christensen, Grünbaum, & Kyllingsbæk*,* 2014 and Kopiske*,* Bruno, Hesse, Schenk, & Franz, 2016). For example, the visual system's tendency to determine absolute size by comparing an object with its surround gives rise to the Ebbinghaus Illusion (see Fig. 2.1), in which a circle surrounded by relatively small circles appears larger than an identical circle surrounded by relatively large circle.

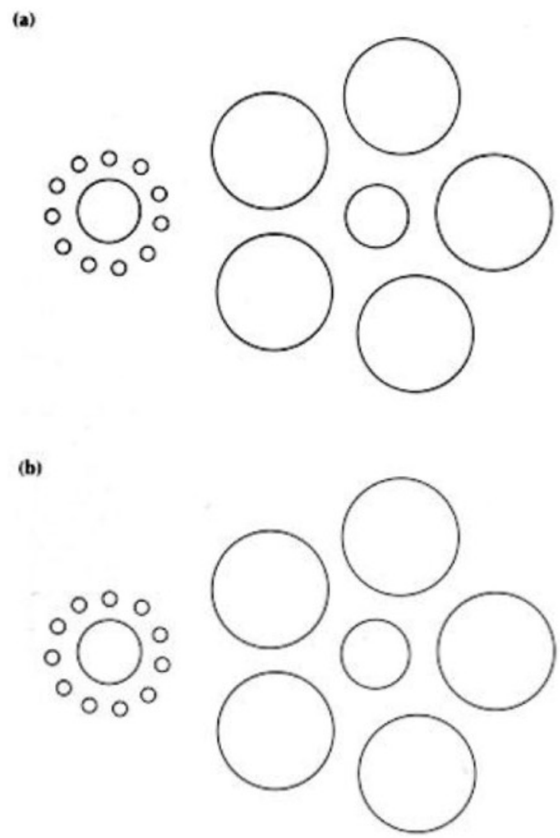

Figure 2.1: (a) The traditional Ebbinghaus illusion. The center circles are the same size, but the left appears larger than the right. (b) A modified illusion in which the center circles appear to be of the same size, but are in fact different.

Aglioti *et al.* (1995) presented subjects with a 3D version of the Ebbinghaus illusion (see Fig. 2.2), such that either (a) the center circles appeared to be of equal size, but were in fact different sizes or (b) the center circles appeared to be of different sizes, but were in fact the same size. They instructed subjects to pick up the right center circle if they believed the center circles to be of equal size and the left center circle if they believed the center circles to be of different size. (The order of left and right were controlled across blocks.) They then measured the maximum grip aperture of the fingers as the subject approached the center disc. What they found was surprising. Holding fixed the actual size of the center circle, the maximum grip aperture was not statistically different between trials in which the center circles appeared to be of identical size and those in which the center circles appeared to be of different size. That is, the Ebbinghaus effect appeared to have minimal influence on the guidance of the subject's grip as they reached for the circle. This suggests that the representation driving this behavior is less sensitive to this comparative size effect than is conscious experience.

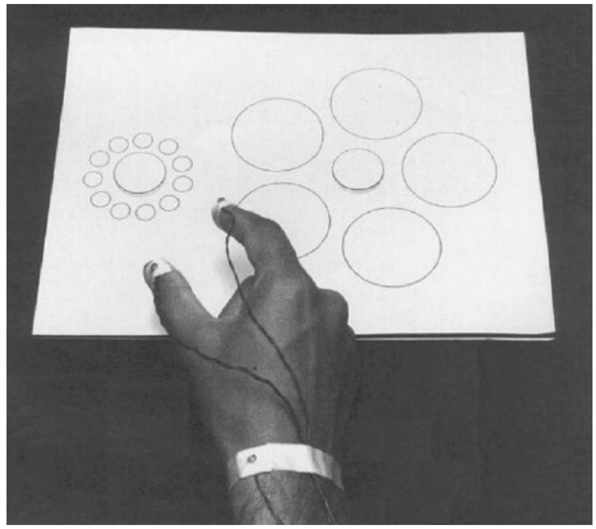

Figure 2.2: The experimental setup from Aglioti *et al.* (1995). The subjects were presented with 3D versions of the traditional (seen here) and modified Ebbinghaus illusions (see Fig. 2.1). In this example, the subject judges that the circles are of different sizes, and reaches for the left center circle.

The fact that visual motor guidance is relatively insensitive to comparative size is direct evidence for representational disunity. That is, it suggests that the representation driving visual motor guidance is less sensitive to the Ebbinghaus illusion than is conscious visual experience. Since these representations have different properties, then, it follows that they're distinct.

Milner and Goodale argue that it makes sense for conscious visual experiences and the representations guiding action to be differentially sensitive to the Ebbinghaus illusion (2006). Visual experiences, they suggest, must locate objects in objective space to enable flexibility to the location and sizes of those objects *diachronically*, as the egocentric relationship with those objects might change; it therefore makes sense that their representations of object size be highly sensitive to contextual information. Guidance representations, by contrast, need only enable the subject to locate the edges of objects in egocentric space as they guide their hand towards them; this task can be less sensitive to contextual cues (2006, section 6.4).

But studies such as these do not tell us whether the behavior is produced by visuomotor modules. It's of course possible that separate modules process spatial information differently in order to realize the kind of representational disunity suggested by Milner and Goodale. But it's also possible that a highly integrated system simply creates multiple representations with different properties for different downstream tasks.

One might conclude that behavioral studies like these only indicate where there is representational disunity, and thus tells us nothing at all about the computational features of the systems which generate that disunity. But that move is too quick. From these results, we know that the mechanism which produces the representation driving motor control must have the ability to relatively accurately determine the sizes of objects, even when presented with potentially misleading contextual information. And that in turn suggests that this mechanism employs a different strategy for determining distal size than does the mechanism which produces conscious representations. We could test different claims about this strategy using behavioral studies similar to this one. If, for example, the mechanism which generates guidance representations estimates sizes using visual direction (that is, the direction of edges in relation to the eye), we should expect those representations to be less sensitive to depth cues (Linton, 2020).

In this way, we can test claims about the computational features of the mechanisms driving behavior. This method is limited, however. Just as we cannot know whether the systems underlying this behavior are modular, we cannot know much about how they interact. Thus, we cannot know whether there is an isolable core mechanism for either of the respective representations. With respect to behavioral studies such as these, then, the modularity hypothesis and the core mechanism hypothesis are on par. Behavioral studies can establish an *explanandum* which could be explained by a wide variety of underlying architectures. The purpose of other sorts of experiments is to help determine further features of that architecture. This way of thinking about the role of behavioral studies in our understanding of the two streams is both straightforward, once articulated, and revisionary, as it suggests a division between the experiments confirming representational disunity from those studying underlying architecture. I'll return to this point in the next section.

A second kind of evidence discussed in the two streams literature comes from lesion studies. Subjects with isolated damage to either the ventral stream or the dorsal stream exhibit deficits in performance, giving us clues about the functional role of the lesioned area. Milner and Goodale's early articulation of their model was motivated by work on patient D.F., whose isolated damage to the lateral occipital area of the ventral stream resulted in a condition known as 'visual form agnosia' a deficit in perceiving visual form. In a series of experiments, Goodale, Milner, Jakobson, and Carey (1991) compared D.F.'s performance on a visual matching task — in which she was asked to visually match the orientation of a card with that of a slot in front of her — with her performance on a visuomotor ‘posting’ task — in which she physically fit the card through the slot. It was found that while her performance on the visual matching task was severely impaired relative to a control subject, reflecting a general deficit in her ability to visually recognize object orientations, her performance on the visuomotor posting task was similar to the control’s (see Fig. 2.3). Milner and Goodale conclude that D.F.’s isolated damage to the ventral stream results in a deficit in visual perception without a corresponding deficit in visuomotor control.

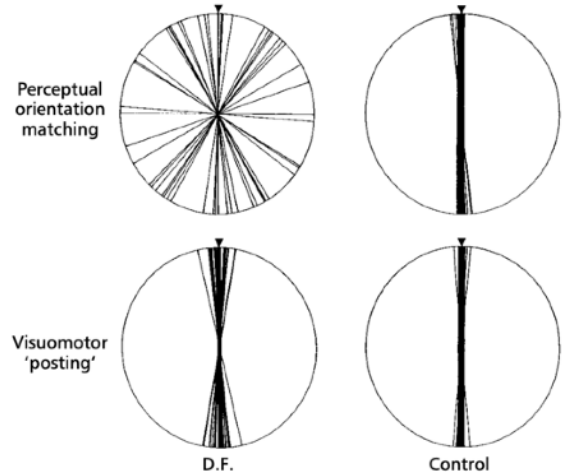

Figure 2.3: Polar plots reflecting the orientation of the cards for D.F. and a control subject in the perceptual matching and visuomotor posting tasks. From Milner and Goodale (2006), data from Goodale *et al.* (1991).

Milner and Goodale compare the case of D.F. with optic ataxics, whose isolated dorsal stream damage results in a deficit in visuomotor control without a corresponding deficit in conscious experience (Milner & Goodale, 2006; for a review, see Andersen, Andersen, Hwang, & Hauschild, 2014). Most optic ataxics have unilateral damage to the posterior parietal lobe, resulting in a deficit in their ability to reach for objects in the contralesional visual field (Riddoch, 1935; Cole, Schutta, & Warrington, 1962; Ratcliff & Davies-Jones, 1972). This deficit is not explained by general damage to spatial processing, since the deficit is only contralesional. Nor is it explained by a deficit general to spatial *perception*, since the deficit is specific to a given effector system. That is, optic ataxics typically perform normally when directing their eyes toward a target (Ratcliff & Davies-Jones, 1972) or, in at least one early case, when performing the task with an unaffected hand (Bálint, 1909). Finally, the deficit is not explained by a general motor disorder, as optic ataxics perform normally when pointing to objects on their body — a task requiring proprioception and touch, but not vision — or when pointing to objects in the ipsilesional visual field (Ratcliff & Davies-Jones, 1972). Thus, the deficit for optic ataxics appears to be a deficit in the representation for contralesional visual motor guidance. Together with patients like D.F., then, optic ataxics appear to establish a double dissociation between visual perception — which is uniquely impacted by ventral stream lesions — and visuomotor control — which is uniquely impacted by dorsal stream lesions. Thus, these results have been taken to strongly support Milner and Goodale's functional characterization of the two streams.

These results however, like cases of double dissociation more generally, do not confirm the existence of separate modules for motor guidance and conscious perception. At least two problems prevent that conclusion. First, visual perception and motor guidance are highly complex and disjunctive phenomena. Thus, it doesn't follow from the fact, for example, that D.F. performs normally on one posting task that she performs normally on all visually guided motor tasks. In fact, she does not. Her performance is substantially worse than neurotypical controls in tasks involving slightly more complex shapes (Goodale *et al.,* 1994; Carey, Harvey, & Milner, 1996). She has trouble gripping objects 'naturally' (i.e. in a way that enables a comfortable grip), even when doing so does not involve semantic recognition of the object (Dijkerman, McIntosh, Schindler, Nijboer, & Milner, 2009; Dijkerman, Milner, & Carey, 1998; McIntosh, Dijkerman, Mon-Williams, & Milner, 2004). Whereas neurotypical subjects use a variety of cues to guide their actions, D.F. appears to rely exclusively on vergence angle (a binocular depth cue) and vertical gaze angle (i.e. visual height) when determining reach distance (Mon-Williams, McIntosh, & Milner, 2001; Mon-Williams, Tresilian, McIntosh, & Milner, 2001). And she employs only binocular disparities and motion parallax to determine object depth for grasping (Dijkerman & Milner, 1998; Dijkerman, Milner, & Carey, 1998, 1999). Interfering with these cues results in her performing substantially worse on motor guidance tasks relative to neurotypical controls (see also Schenk & McIntosh, 2010). After having reviewed the now substantial evidence for ventral stream influence on visuomotor control, Briscoe and Schwenkler (2015) suggest that dorsal stream processes directly recruit ventral stream representation when engaged in complex or novel guidance tasks (for similar claims of direct ventral stream influence on dorsal stream processes, see Schenk & McIntosh, 2010 and McIntosh & Lashley, 2008). Thus, even if there is a double dissociation between Milner and Goodale's perceptual matching and visuomotor posting tasks, there can be (and is) substantial interaction between the systems underlying these behaviors.

Suppose, however, that the double dissociation did hold in lesion studies across the many multitudes of task-types that make up visual perception and visuomotor guidance. A second problem is that the inference from double dissociation to modularity is strong only when the relevant lesions are both complete and precise. By 'complete,' I mean that a lesion knocks out the entire region attributed to a particular task. Thus, a dorsal stream lesion is complete in this sense only if it affects the entirety of the dorsal stream. By 'precise,' I mean that a lesion renders inoperable only the target region. Like lesion studies generally, the cases studied in this literature are certainly neither complete nor fully precise (Schenk & McIntosh, 2010). And when a lesion is incomplete or imprecise, the possibility remains that an observed effect is underwritten either by the unknocked-out part of the target region or by the knocked-out part of the non-target region. Thus, for example, since D.F.'s lesion knocks out only an isolated part of the ventral stream, it remains possible that other parts of the ventral stream help to determine her visuomotor performance.

Both problems turn on the fact that modularity entails the *irrelevance* of *non-target regions* to a task, something lesions studies do not directly test. By contrast, a region's being a core mechanism entails only the *central relevance* of that *target region* to the task, which lesion studies do directly test. Thus, it's not a challenge to a core mechanism claim that different versions of a task recruit different regions of cortex (even regions outside the core mechanism), such that lesions on those areas impact performance. What matters is that lesions to the core mechanism are associated with more far-reaching deficits in the target task. And lesion studies confirm precisely that: while D.F.'s ventral stream lesion is associated, contra Milner and Goodale, with specific deficits in motor guidance as described above, dorsal stream lesions (such as those of optic ataxics) are associated with a broad ranging inability to perform visuomotor guidance tasks (Andersen *et al.*, 2014). And the converse case is even clearer: while dorsal stream lesions are associated with certain specific impacts on perceptual experience (Schenk & McIntosh, 2010), lesions to the ventral stream are associated with profound impacts, including complete cortical blindness (Milner & Goodale, 2006).

And while the extent of lesions remains relevant to the evidential support corresponding behavior gives to a core mechanism claim, the requirements are not as strict. It's enough to establish positive relevance of a region to a task that that lesion be *precise*, i.e. not extend beyond the boundaries of the target region. If so, then since a deficit in performance can reasonably be attributed to the lesion, one can reasonably infer the relevance of that region to the task. That is, since the region includes the entire lesion, we can reasonably say that at least some part of the region is associated with the task. Thus, it's not important to establish positive relevance of a region to a task, that a lesion be *complete*, i.e. to encompass the entire target region. This is crucial, as while none of the relevant lesion cases discussed in the literature are complete, many of them are relatively precise (Schenk & McIntosh, 2010). Thus, lesion studies remain a useful indicator of core mechanisms.

A final source of evidence for the two streams can be grouped, not by experimental methodology as in the last two examples, but by the kind of analysis through which experimental evidence is interpreted. In what I'll call computational analysis, researchers employ particular theories about the functions of the streams to make predictions about the computational properties of those streams. These predictions are then tested using a variety of methods, including imaging studies (such as FMRI), knock-out and lesion studies, and so forth.

Milner and Goodale (2006) argue that visuomotor guidance requires representational systems which have a relatively high temporal resolution, update frequently, are relatively coarse-grained, present spatial information egocentrically, and so forth. Perceptual inputs to cognition, by contrast, must have a high spatial resolution, enable long storage duration, locate objects in allocentric space, etc. These claims, along with their claim that the dorsal stream underwrites visuomotor guidance while the ventral stream underwrites conscious perceptions, entail substantive predictions about the nature of processing within the two streams. And these predictions contrast sharply with competitor models, such as Ungerleider and Mishkin's (1982) 'what vs. where' account. Where the latter predicts that the ventral stream is unique in processing information related to object identity while the dorsal stream is unique in processing spatial information, Milner and Goodale's model predicts that each stream should separately process both object identity and spatial information, though perhaps in different ways.

The predictions of Milner and Goodale's model are largely borne out by the evidence. The dorsal stream takes its inputs primarily from fast, but coarse-grained magno-cellular pathways in the LGN, while the ventral stream takes its inputs primarily from slow, but fine-grained parvo-cellular pathways (Milner & Goodale, 2006, section 2.3). These input features, along with other processing features within the respective streams, appear to be responsible for relatively quick-and-dirty outputs from the dorsal stream as compared to relatively slow-and-detailed outputs from the ventral stream. There is direct evidence for spatial processing within each stream, with retinotopic and egocentric organization persisting throughout dorsal stream processing and relatively retino-neutral and allocentric organization (Milner & Goodale, 2006, section 4.1). Both the ventral and dorsal streams appear to process information related to object identities, with the former processing fine-grained information required for categorization and inference and the latter processing relatively coarse-grained information required for navigation and manipulation (Milner & Goodale, 2006, section 2.5.4). Finally, as predicted by Milner and Goodale's account, but not Ungerleider and Mishkin's, the dorsal stream's outputs are predominantly to motor cortices, while the ventral stream's outputs are to prefrontal regions (Milner & Goodale, 2006, sections 2.5 & 2.6).

As we'll see in the next section, computational analysis is, in my view, the core explanatory work to be done by an account of the visual streams. In contrast to lesion studies, computational analysis establishes not just *that* particular regions of cortex are predominantly responsible for the performance of certain person-level roles, but *how* they perform those roles. It is also therefore the best evidence that the modules are core mechanisms: the received view in the literature, even among Milner and Goodale's opponents, is that the central computations required to generate the characteristic representations of the respective streams are performed within those streams (e.g. Schenk & McIntosh, 2010). This is again not to say that those computations are isolated from goings-on outside of the streams; it is merely to say that the anatomical division between the ventral and dorsal streams appears to reflect a *relative* specialization for motor control and conscious perception.

By contrast, the claim that the streams are visuomotor modules is not well supported by computational considerations. It's not just that, as reviewed above, there is evidence that subjects with lesions to one stream have deficits in their performance on the other stream's characteristic task, such as D.F.'s ventral stream damage giving rise to a deficit in certain kinds of motor tasks. We now have evidence for the *particular way* in which representations in one stream appear to influence computations in the other stream. Thus, for example, there is evidence that the dorsal stream recruits ventral stream representations to enable sensitivity to allocentric features (Schenk & McIntosh, 2010), pictorial depth (Schenk & McIntosh, 2010), and object identity (van Polanen & Davare, 2015) when engaging in tasks requiring complex fine-tuning of grasp. Such evidence directly undermines the claim that the streams are visuomotor modules for different purposes. But it simultaneously informs our understanding of the mechanisms responsible for the generation of conscious experiences and guidance representations respectively. It suggests, for example, that the dorsal stream may have to transform ventral stream object representations for the purposes of motor guidance (van Polanen & Davare 2015). Thus, even as such evidence undermines a conception of the visual streams as modules, it enriches our understanding of the core mechanisms operative in the performance of their characteristic tasks.

2.3 The Visual Streams in the Study of Perception

I began this chapter by noting that, if the argument from Chapter 1 is correct, the central implications of Milner and Goodale's model derive, not from their commitment to isolated perceptual processing streams, but from their commitment to representational disunity. But representational disunity can be realized by a variety of underlying architectures, including by multiple modules (as Milner and Goodale suggest), by a more centralized perceptual processing system (as their opponents typically maintain), or by separate core mechanisms (as I've now defended). The challenge that remains, then, is to explain why these different architectures matter, given that none is required for representational disunity and thus for the central implications. In this section, I first articulate the role that visual streams — understood as core mechanisms — can play in the study of perception. I then explore the implications of having them perform this role. Visual streams as core mechanisms, I claim, have earned their place in the ontology of the perceiving mind.

In articulating an account of mechanisms, I noted that mechanistic identity is relative to a previously determined explanandum phenomenon. In the case of the visual streams, that phenomenon is the existence and features of disunified visual representations. The mechanism or set of mechanisms which explain this feature must explain, first, how the visual system transforms retinal stimulation into the set of representations that drive person-level phenomena such as program selection and guidance (the existence question), and, second, how those transformations give rise to the characteristic features of those representations (the features question). It follows that on the mechanistic approach, visual streams are explanatorily secondary to representational disunity. That is, visual streams get their identity conditions and explanatory force *in virtue of* the downstream representations they characteristically produce.

This picture of the role of the visual streams is in sharp contrast with the modularity approach. Modules do not get their identity conditions from the phenomena they explain. Instead, modules have their identity conditions in virtue of the set of features described at the beginning of Section 2.1. The most central of those conditions — informational encapsulation, inaccessibility, domain specificity, and so forth — are *computational* features of mental units. It follows that one can discover modules — independently of a determination of their characteristic outputs — by searching for mental units whose processes operate only over local information stores, which are inaccessible to processes outside, and so on. This independence is reflected in the history of the two streams literature: the existence of separate pathways of visual perception were posited long before researchers had precise theories about their function, and theories about their characteristic output have widely differed.

Thus, on the mechanistic approach, one first discovers representations with certain features, and only then determines the mechanisms responsible for those representations and their features. But, on the modularity approach, one can — and in the case of the two streams literature, researchers thought they did — discover modules first, and only then discover the nature of the representations those modules characteristically produce.

This difference in the order of discovery also impacts standard assumptions about the explananda and explanans operative in the positing of visual streams. As discussed in Chapter 1, it is typically taken for granted that the central implications of the two streams account follow from the positing of separate visuomotor mechanisms with different functional roles. Thus, it is assumed that those implications follow only if the visual streams satisfy the conditions for modularity. But that assumption is false: the implications follow from representational disunity, not the strictly stronger modularity account.

Thus, by contrast, on the picture described above, explanations come in two stages. In the first stage, disunified representations and their different features explain certain features of behavior. In the second stage, a mental architecture (such as a core mechanism) explains the existence and features of those disunified representations. This two-stage framework better tracks the dependencies between disunified representations, underlying architecture, and the implications of the two.

The different frameworks also posit different explanatory roles for the three kinds of studies discussed in the last section. As discussed, Milner and Goodale appear to treat behavioral studies, lesion studies, and computational analyses as all jointly confirming the existence of separate visuomotor modules for different tasks. But I argued that the central implication of the behavioral evidence is representational disunity, which can be realized by a variety of underlying architectures. While lesion studies can help discover mental architectures, they're better suited for discovering causal relationships between a region and phenomenon, not the informational isolability of that region. To discover the broader sets of interactions between a region and other regions, one must perform broader computational analyses on mental systems.

In contrast to the mechanistic approach, the two-stage framework separates the discovery of disunified representations from the discovery of their underlying architecture. It can thus better capture the evidential role of the three kinds of study. That is, on my account, the primary role of behavioral evidence is to help discover disunified representations. The role of lesion studies is to determine the causal relevance of neural regions to those representations (and thus to discover which regions are components of the mechanism for those representations). And the role of computational analyses is to discover how different regions interact in order to produce those representations (and thus to discover whether those representations are constituted by unified systems, modules, or separable core mechanisms).

Thus, I think the two-stage project employed in the mechanistic approach has clear advantages over the project typified in the modularist approach. I'll finish my discussion by highlighting three sorts of meaningful explanatory work that is accomplished by the mechanistic approach. This work, I believe, enables the mechanistic notion of streams to earn its place in our mental ontology.

First, positing core mechanisms enables us to explain and predict the existence and features of disunified representations. Core mechanisms explain the existence and features of disunified representation because, as described above, to discover a mechanism is to discover the entities and activities which causally mediate a phenomenon's inputs and outputs. In this case, the relevant outputs are representations with certain features (including functional features) and the inputs are paradigmatically retinal stimulation. Thus, the relevant mechanism is the set of causal factors which explain how retinal stimulation is transformed into those representations. This set of causes explains both how the representations come to exist and why they have the features they do. Thus, for example, if — as Milner and Goodale contend — dorsal stream processes operate over egocentrically formatted representations, then this fact can explain the spatial format of action-guidance representations (which, Milner and Goodale also contend, are themselves egocentrically formatted). And in addition to explaining the features of disunified representations, knowing the mechanism which causes those representations can enable us to predict some of their features. Thus, to reverse the preceding case, if we antecedently knew that dorsal stream processes operate over egocentrically formatted representations, we could justifiably predict that guidance representations are also egocentrically formatted.

The preceding points are true of mechanistic explanation generally. What additional explanatory work is achieved by positing a neurally individuable *core* mechanism for disunified representations? Recall that a core mechanism is the local mechanism which explains the central features of a target phenomenon. Thus, when the core mechanisms for different representations are neurally individuable, we know that the central features of those representations are determined by separate processing units. This does not entail the positive claim that the representation and their causes can be studied in isolation, as might be entailed by modularity. But it does entail the negative claim that we cannot make direct inferences from one representation (or its causes) to the other representation (or its causes). Thus, if disunified representations are determined by neurally individuable core mechanisms, then we must be cautious when drawing inferences about, for example, visual perception generally.

Second, positing visual streams as core mechanisms has important architectural implications. Specifically, the notion of visual streams developed above retains the idea that the visual system divides and conquers, employing separate perceptual processing streams for separate functions. You might doubt this implication on the grounds that the requirements of core mechanisms are weaker, and thus more architecturally neutral, than those of modularity. As I pointed out above, the claim that the visual streams constitute neurally isolable core mechanisms is neutral between with the claim that those streams are *in fact* modules (as the latter is just a kind of core mechanism) and the claim that the streams are part of a highly interconnected processing system (as in fact seems likely of the visual streams). Given this neutrality, one might doubt the architectural implications of positing core mechanisms.

But this places too high a bar on the discovery of mental architecture. As I suggested when developing the mechanistic account, neural processes are paradigmatically global, in that any given mental phenomenon is typically underwritten by many sets of overlapping systems which interact in endlessly complex ways to bring about a phenomenon. Even if there are mental modules, it's an important aspect of neuroscientific work that we break down highly integrated systems into relatively local mechanisms which realize core functions. Thus, the discovery of neurally isolable core mechanisms for separate roles is an important part of the discovery of mental architecture. Put another way, to say, for example, that the dorsal stream is the core mechanism for the generation of guidance representations is on par with saying that the hippocampus constitutes the core mechanism for episodic memory. Both mechanisms are of course widely integrated with systems outside. But they remain important architectural findings in the study of their respective phenomena.

Finally, in addition to allowing us to explain and predict features of disunified representations in virtue of features of core mechanisms, knowing that a set of processes constitute a core mechanism can enable us to explain and predict features of its internal computations. That is, for example, Milner and Goodale argue that program guidance is assisted by processes that are quick, coarse grained, have an egocentric format, and so on. This contention, along with the claim that the dorsal stream constitutes the core mechanism for guidance representations, entails positive predictions about the goings on within the dorsal stream: namely, we can expect dorsal stream processes to have these features. In this respect, the positing of neurally-isolable core mechanisms does similar work to Milner and Goodale's positing of separate visuomotor modules: it allows us to make predictions about the *different* computational features of different streams. This in turn helps guide researchers.

References

Aglioti, S., DeSouza, J. F. X., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, *5*, 679–685.

Andersen, R. A., Andersen, K. N., Hwang, E. J., & Hauschild, M. (2014). Optic Ataxia: From Balint's Syndrome to the Parietal Reach Region. *Neuron*, *81*(5), 967–983. https://doi.org/10.1016/j.neuron.2014.02.025

Bálint, R.D. (1909). Seelenlähmung des "Schauens", optische Ataxie, räumliche Störung der Aufmerksamkeit. pp. 67–81. *European Neurology*, *25*(1), 67-81.

Binkofski, F., Buxbaum, L.J., 2013. Two action systems in the human brain. Brain Lang. 127, 222–229. http://dx.doi.org/10.1016/j.bandl.2012.07.007.

Briscoe, R., & Schwenkler, J. (2015). Conscious Vision in Action. *Cognitive Science*, *39*(7), 1435–1467. https://doi.org/10.1111/cogs.12226

Carey, D. P., Harvey, M., & Milner, A. D. (1996). Visuomo- tor sensitivity for shape and orientation in a patient with visual form agnosia. *Neuropsychologia*, *34*(5), 329–337.

Christiansen, J. H., Christensen, J., Grünbaum, T., & Kyllingsbæk, S. (2014). A Common Representation of Spatial Features Drives Action and Perception: Grasping and Judging Object Features within Trials. *PLoS ONE*, *9*(5), e94744. https://doi.org/10.1371/journal.pone.0094744

Cole, M.C., Schutta, H.S., and Warrington, E.K. (1962). Visual disorientation in homonymous half-fields. *Neurology*, 12, 257–63.

Craver, C. F. (2007). *Explaining the Brain*. Oxford University Press.

Dijkerman, H. C., McIntosh, R. D., Schindler, I., Nijboer, T. C. W., & Milner, A. D. (2009). Choosing between alternative wrist postures: Action planning needs percep- tion. *Neuropsychologia*, *47*, 1476–1482.

Dijkerman, H. C., & Milner, A. D. (1998). The perception and prehension of objects oriented in the depth plane. II. Dissociated orientation functions in normal subjects. Experimental Brain Research, 118, 408–414.

Dijkerman, H. C., Milner, A. D., & Carey, D. P. (1998). Grasping spatial relationships: Failure to demonstrate allocentric visual coding in a patient with visual form agnosia. *Consciousness and Cognition*, *7*, 424–437.

Dijkerman, H. C., Milner, A. D., & Carey, D. P. (1999). Prehension of objects oriented in depth: Motion parallax restores performance of a visual form agnosic when bin- ocular vision is unavailable. *Neuropsychologia*, *37*, 1505–1510.

Fodor, J. (1983). The modularity of mind: An essay on faculty psychology. MIT press.

Goodale, M. A., Jakobson, L. S., Milner, A. D., Perrett, D. I., Benson, P. J., & Hietanen, J. K. (1994). The nature and limits of orientation and pattern processing support- ing visuomotor control in a visual form agnosic. *Journal of Cognitive Neuroscience*, *6*, 46–56.

Goodale, M. A., & Milner, A. D. (2018). Two visual pathways – Where have they taken us and where will they lead in future? *Cortex*, *98*, 283–292. https://doi.org/10.1016/j.cortex.2017.12.002

Goodale, M.A., Milner, A.D., Jakobson, L.S., and Carey, D.P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, 349, 154–6.

Ingle, D. (1973). Two visual systems in the frog. *Science*, 181, 1053–5.

Knudsen, E. I. (2007). Fundamental components of attention. *Annu. Rev. Neurosci.*, *30*, 57-78.

Kopiske, K. K., Bruno, N., Hesse, C., Schenk, T., & Franz, V. H. (2016). The functional subdivision of the visual brain: Is there a real illusion effect on action? A multi-lab replication study. *Cortex*, *79*, 130–152. https://doi.org/10.1016/j.cortex.2016.03.020

Linton, P. (2020). *Conflicting shape percepts explained by perception cognition distinction.* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/27wjr

Marr, D. (1982). *Vision*. Freeman, San Francisco.

McIntosh, R. D., Dijkerman, H. C., Mon-Williams, M., & Milner, A. D. (2004). Grasping what is graspable: Evi- dence from visual form agnosia. *Cortex*, *40*, 695–702.

McIntosh, R. D., & Lashley, G. (2008). Matching boxes: Familiar size influences action programming. *Neuropsy- chologia*, *46*, 2441–2444.

Michel, F., & Henaff, M. A. (2004). Seeing without the occipito-parietal cortex: Simultagnosia as a shrinkage of the attentional visual field. *Behavioural Neurology*, *15*, 3–13.

Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action* (2nd ed.). Oxford, UK: Oxford University Press.

Mole, C. (2009). Illusions, Demonstratives, and the Zombie Action Hypothesis. *Mind*, *118*(472), 995–1011. https://doi.org/10.1093/mind/fzp109

Mon-Williams, M., McIntosh, R. D., & Milner, A. D. (2001). Vertical gaze angle as a distance cue for pro- gramming reaching: Insights from visual form agnosia II (of III). *Experimental Brain Research*, *139*, 137–142.

Mon-Williams, M., Tresilian, J. R., McIntosh, R. D., & Milner, A. D. (2001). Monocular and binocular distance cues: Insights from visual form agnosia I (of III). *Exper- imental Brain Research*, *139*, 127–136.

Perenin, M. T., & Vighetto, A. (1988). Optic ataxia: A spe- cific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain*, *111*, 643–674.

Pisella, L., Sergio, L., Blangero, A., Torchin, H., Vighetto, A. & Rossetti, Y. (2009). Optic ataxia and the function of the dorsal stream: Contribution to perception and action. *Neuropsychologia.* doi:10.1016/j.neuropsychologia. 2009.06.020

Prychitko, E. (2019). The causal situationist account of constitutive relevance. *Synthese*, 1-15.

Ratcliff, G. and Davies-Jones, G.A.B. (1972). Defective visual localization in focal brain wounds. *Brain*, 95, 49–60.

Riddoch, G. (1935). Visual disorientation in homonymous half-fields. *Brain*, 58, 376–82.

Schenk, T., & McIntosh, R. D. (2010). Do we have independent visual streams for perception and action? *Cognitive Neuroscience*, *1*(1), 52–62. https://doi.org/10.1080/17588920903388950

Striemer, C., Blangero, A., Rossetti, Y., Boisson, D., Rode, G., Vighetto, A., et al. (2007). Deficits in peripheral visual attention in patients with optic ataxia. *Neuroreport*, *18*, 1171–1175.

Striemer, C. L., Chapman, C. S., & Goodale, M. A. (2009). “Real-time” obstacle avoidance in the absence of prim- ary visual cortex. Proceedings of the National Academy of Sciences of the United States of America, 106, 15996–16001.

Striemer, C., Locklin, J., Blangero, A., Rossetti, Y., Pisella, L., & Danckert, J. (2009). Attention for action? Examin- ing the link between attention and visuomotor control deficits in a patient with optic ataxia. *Neuropsychologia*, *47*, 1491–1499.

Tulving, E. (1983). *Elements of episodic memory*. Clarendon.

Ungerleider, L.G. and Mishkin, M. (1982). Two cortical visual systems. In *Analysis of visual behavior*, (ed. D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield), pp. 549–586. MIT Press, Cambridge, MA.

van Polanen, V., & Davare, M. (2015). Interactions between dorsal and ventral streams for controlling skilled grasp. *Neuropsychologia*, *79*, 186–191. https://doi.org/10.1016/j.neuropsychologia.2015.07.010

Chapter 3: The Function of Functional Division

3.1 Introduction

One of the features that makes cognitive science a science is its insistence on decomposition. The mind is not a uniform input-output machine, but rather a collection of discrete states, processes, and systems which collectively determine our mental lives. Therefore, one of the central tasks of cognitive science is to describe the functions of discrete mental units. These functional descriptions can come in two broad forms. *Local* functional descriptions characterize the central tasks of a mental unit and the internal operations which realize that performance. *Global* functional descriptions characterize, of a collection of mental units, how they interact in order to realize mental phenomena.

Local and global functional descriptions are distinct, but tightly related, such that each is incomplete without the other. Failing to attend to both local and global questions leads us to miss things. Loosely speaking, the local function of a corner store is to sell food and other home goods. But when attending to the global function of home goods stores more generally, we see that corner stores are one of many kinds of stores which employ different strategies to deliver home goods to customers. This explains why corner stores carry Twinkies, but not produce, and cleaning solution, but not mops. Thus, attending to the global function of home goods stores causes us to refine our characterization of the local function of corner stores.

Some functional questions can only be asked *globally*, that is, of suites of discrete units. We might ask, for example, what broader role the division between corner stores, groceries, and hardware stores plays in the sale of home goods. Why not just have a single store that carries everything? The answer is (approximately) that the distinction between groceries and hardware stores allows for the concentration of resources and expertise according to distinct needs, while the distinction between these two and corner stores broadens both speed and access. This answer forces us to further refine our characterization of the local function of corner stores.

Like the home goods market, the visual system is subdivided into discrete units with distinct local functions. Much of this division is what we might call sub-personal: there's a visual unit for motion detection, edge and object detection, color, and so forth. These units coordinate in the production of broadly unified visual experiences, each containing components processed in these various units. But the visual system is also divided according to *person-level* local functions. On the dominant accounts, the cortical visual system is divided into a ventral stream processing 'vision for perception' and a dorsal stream processing 'vision for action' (Milner & Goodale, 2006), while the subcortical visual system is divided into a 'low road' for rapid threat detection and response and a 'high road' which outputs to the cortical system (Tamietto & De Gelder, 2010).

This division in person-level local functions raises a higher-order *global* question: what is the function of this functional division? Why does the visual system divide and conquer rather than unite and build? On the dominant account, designed to explain the functional division between the ventral and dorsal streams, the division between visual streams is similar to the division between groceries and hardware stores: it allows for the concentration of resources and expertise according to distinct needs. That is, perception and motor control place conflicting demands on the computational systems that determine them, and adequately meeting those demands is best accomplished with a divide and conquer strategy (Milner & Goodale, 2006). I'll call this view "task-based division." In their critique of this model, Briscoe and Schwenkler (2015) suggest an alternative, apparently deflationary, account according to which the division is more like that between groceries and corner stores: both the ventral and dorsal streams are ultimately engaged in action guidance, but the ventral stream is slow and deliberate while the dorsal stream is quick and automatic. I'll call this view "computational division."

My aim in this chapter is to defend a new account of the function of functional division which is intermediate between task-based division and computational division. The function of functional division is to allow, for each task, a division between task-specific and task-general processing strategies, optimizing performance. In the next two sections, I consider and ultimately reject task-based and computational division, respectively. Because multiple streams participate in the performance of their proposed person-level tasks, the function of functional division is not simply, as Milner and Goodale propose, to assist in the performance of downstream tasks with conflicting demands (Section 3.2). And because some streams are largely dedicated to the performance of particular downstream tasks, the function of functional division is not simply, as Briscoe and Schwenkler suggest, to realize different computational strategies in response to the same sorts of tasks (Section 3.3).

In Section 3.4, I propose my new view. I argue that, to realize their computational advantages, automatic processing systems organize around particular person-level tasks while deliberate ones are task neutral. Thus, realizing the benefits of multiple computational strategies requires a fundamental division between what I call 'task-coupled' and 'task-decoupled' processing systems. Such a division explains the empirical findings that challenge task-based and computational division.

3.2 Task-based Division

In this section I'll present Milner and Goodale's (2006) account of the function of functional division, which I'll call 'task-based division', and argue that it is inconsistent with the empirical evidence.

3.2.1. Local Functions

Neuroscientists have long posited a division in the cortical visual system between the ventral and dorsal streams. I mentioned above that on Milner and Goodale's (2006) account, the ventral stream has the function of processing vision for perception while the dorsal stream has the function of processing vision for action. By 'vision for perception,' they mean the set of representations (and upstream processing) which are operative in (largely conscious) cognitive activity, such as conscious visual report. By 'vision for action,' they mean the set of representations and processes which are operative in in-the-moment visual motor guidance, such as the visual guidance of the hand when reaching for an object.

To help explain the broader functional role of such a division, they draw an analogy to the guidance of a hypothetical martian rover (Goodale & Milner, 2004). The rover's job is to navigate the martian terrain and collect samples. To accomplish these goals, the rover is fitted with communication equipment which allows remote control of its overall behavior from Earth. To determine the rover's action plan, mission control determines what the rover's sensors tell us about the rover's nearby environment and sends an action program describing what it should do next. But since there is substantial delay between Earth and Mars, it would be inadvisable for the rover to simply execute precise movements on the basis of mission control's action program.

Rather, the rover should update the program in response to in-the-moment conditions. Thus, the rover is fitted with an additional onboard system which independently assesses the scene and makes fine-grained adjustments to the action program. If a rock blows in the way, it can navigate around it; when picking up samples, it can use its sensors in-the-moment to guide its limbs toward the target; and so forth. The overall behavior of the rover is thus controlled by two systems, a smart but slow controller system which *selects* action programs for the rover, and a crude but quick onboard system which *guides* the fine-grained execution of those programs.

Milner and Goodale propose that a similar functional division exists in the cortical visual system. The ventral stream outputs to a broad-functioning cognitive system, enabling smart but slow *selection* of action programs; and the dorsal stream outputs directly to action systems to enable dumb but quick *guidance* of those programs. If, for example, I aim to catch a ball hit into the outfield, ventral stream information is employed in determining what to do and roughly how to do it. I assess the trajectory of the ball, see that it is headed toward my part of the field, and make a rough estimate of where I need to move my body and arm to intercept it. Thus, the ventral stream and downstream cognitive activity are analogous to the Martian rover's controller system. But the action program selected by cognition is itself vague; roughly, cognition's output is a set of conditional motor commands. The dorsal stream's role, analogous to the Martian rover's onboard system, is to supply 'online' the antecedents of those conditionals. As the ball arcs, I move my body in order to keep a constant angle between the horizon and the ball, ensuring that I arrive at the bottom of the arc, and so forth. The ventral stream and cognition determine what to do and the dorsal stream rapidly updates action systems about how to do it. In keeping with the rover analogy, I'll refer to these two tasks as program 'selection' and program 'guidance'.

In what follows, I'll assume that these local functional claims are largely correct; that is, that the division between the ventral and dorsal streams roughly corresponds to a division in processing between vision for perception and vision for action. My defense of that model is not complete, however, and it's worth clarifying my commitments by responding to two kinds of evidence which have been raised against the model. Specifically, on Milner and Goodale's preferred interpretation, the ventral and dorsal streams constitute isolated perceptual processing systems, each with a proprietary function. But each of these claims is challenged in the empirical literature. First, evidence from neuroanatomical studies reveal substantial connectivity between the ventral and dorsal streams (van Polanen & Davare, 2015; Cloutman, 2013; Grafton, 2010) which likely results in computation-level interaction between them (Schenk & McIntosh, 2010). Such evidence undermines Milner and Goodale's contention that the ventral and dorsal streams constitute isolated perceptual processing systems. Second, Milner and Goodale's contention that each stream performs a proprietary function (of program selection or guidance respectively) is undermined by evidence for direct dorsal stream influence on conscious experience (Michel & Henaff, 2004; Perenin & Vighetto, 1988; Pisella *et al.*, 2009; Striemer *et al.*, 2007, Striemer, Chapman, & Goodale, 2009; Striemer *et al.*, 2009; van Polanen & Davare, 2015) and for direct ventral stream involvement in online motor control (Goodale et al., 1994; McIntosh & Lashley, 2008; Schenk & McIntosh, 2010; Briscoe & Schwenkler, 2015).

As argued in Chapter 1, neither interconnectivity between the streams nor failure of proprietary function undermine the central tenet of Milner and Goodale's proposal, which is that dorsal stream representations exhibit direct control on online motor activity, without further cognitive intermediaries. This claim is both novel and surprising, since it undermines a picture of the perceiving mind according to which perceptual systems function as inputs to cognition, which then does the heavy lifting in determining our actions. More generally, as argued in Chapter 2, the claim of broad functional division between the streams—that is, that the dorsal stream is largely organized around program guidance while the ventral stream is largely organized around broader cognitive guidance, and that the relevant representations driving those activities are located within those respective streams—is consistent with cross-talk between the two streams, and with partial overlap in their functional roles. As Schenk and McIntosh (2010) put it in their review of evidence against fully isolated perceptual processing streams: "the perception-action model captures some broad patterns of functional localization, but… the specializations of the two streams are relative, not absolute."

Once we've posited multiple streams of visual processing, it's reasonable to ask: how many streams are there? Some researchers building connectivity models of the visual system have proposed that the cortical visual system's topology divides beyond the two visual streams (e.g. Haak & Beckmann 2018). But in Chapter 2, I argued that we should individuate perceptual streams according to the person-level functions their representations primarily subserve. Thus, since not all divisions in processing necessarily entail person-level functional differences, such topological studies do not suffice to establish further visual streams.

Researchers studying the *subcortical* visual system, however, have proposed a division in processing between 'low' and 'high' roads of visual processing. The high road functions to transport visual information from the retina to the visual cortices (and thus, eventually, along the two visual streams). It thus lacks a distinctive person-level function. But the low-road—which is constituted by direct projections from the thalamus (an early 'way station' in visual processing) to the early amygdala (which is associated with fear processing—is thought to directly trigger rapid 'fight or flight' responses to threatening stimuli. These responses include broader anticipatory responses — such as up regulating cortisol production (van Honk *et al.*, 1998) — but also include overt behaviors — such as freezing or fleeing (Hamm *et al.*, 2003). The latter is an example of program *selection,* a paradigm person-level task. Thus, it's plausible that the low-road constitutes a third visual stream. As we'll see, one of my objections to task-based division relies in part on the widely accepted claim that the low-road exists and is independently responsible for triggering a select class of 'fight or flight' action programs.

3.2.2. Global Functions

While the idea of a division between the ventral and dorsal streams is old, the role this division plays in visual processing has long been debated. According to Mishkin and Ungerleider's (1982) earlier account, the ventral stream has the function of determining the 'qualities' of objects in the visual scene while the dorsal stream has the function of determining the locations of those objects. In slogan form, the ventral stream is the 'what' pathway while the dorsal stream is the 'where' pathway. Milner and Goodale's (2006) account of the division—on which the division between the ventral and dorsal stream functions to accomplish separate downstream tasks, rather than process different distal features—has largely supplanted Mishkin and Ungerleider's. Key evidence against the latter includes evidence for separate object and spatial representations present in each stream (Milner & Goodale, 2006). If each stream relies on its own object and spatial representations, then the division between the ventral and dorsal streams cannot be explained by a division between 'what' and 'where' processing.

I bring up this earlier model to make a different point, however: In providing an account of the local function of the two streams, Ungerleider and Mishkin's account makes harder the global question of what purpose such functional division has in the first place. Other divisions in the processing of distal features—such as between motion and color processing in early vision—are common throughout the visual system. In those cases, however, distal features are processed separately but in close spatial proximity. This proximity enables rapid cross-talk when determining the distal features with related contents. And it makes sense given the eventual aim of producing a unified representation of the distal scene. But both of these features are also true of 'what' and 'where' information, since, for example, information about spatial relationships can inform our understanding of object identity and since our visual experiences clearly contain unified representations of objects and their spatial relationship. Why, then, should 'what' and 'where' information be processed in such disparate regions of cortex when other divisions are not? No clear answer is forthcoming. The failure of the Ungerleider and Mishkin model thus highlights the importance of attending simultaneously to both local and global functional questions.

Milner and Goodale (2006) propose a task-based division in processing between the ventral and dorsal streams. The ventral stream outputs to a broad-functioning cognitive system, enabling smart but slow selection of action programs; and the dorsal stream outputs directly to action systems to enable dumb but quick guidance of those programs. This is their account of the *local* function of the two visual streams. Unlike Ungerleider and Mishkin, Milner and Goodale also attend to global functional questions about the broader function of this functional division. They argue that flexible program selection and rapid program guidance place conflicting demands on the computational demands that underly them. Thus, adequate performance on both tasks requires a ‘divide and conquer’ strategy.

They highlight several demand conflicts between selection and guidance (Milner & Goodale, 2006). First, rapid program guidance requires high temporal resolution while flexible program selection requires high spatial resolution. To catch a falling pen, for example, the representations guiding my hand toward the target must be rapidly updated in the moment, allowing me to constantly update the trajectory of my arm as the pen falls. By contrast, determining the identity of an acquaintance (an aspect of program selection) requires attention to fine-grained details of their face. But speed and detail place conflicting demands on computational systems: the faster the processing, the less fine-grained detail can be sussed out; the more detail, the longer it takes to process.

Second, flexible program selection requires long storage duration, while rapid program guidance requires rapid updating. To decide which button to press in the elevator, for example, I need to know which button it was appropriate to press last time. Thus, the system responsible for selection must have access to long-stored contents. By contrast, when reaching for the falling pen, the precise angle between my arm and the pen matters a great deal in the moment, but not at all a few milliseconds later. The representations guiding my hand, therefore, must update in-the-moment. Since storage requires resources, and since rewriting stored representations takes time and energy, long storage duration and rapid updating place conflicting demands on computational systems.

Third, flexible selection is best performed from within an allocentric (object-centered) spatial frame, while rapid guidance is best performed from within an egocentric (agent-centered) spatial frame. To decide where to look for my keys, for example, what matters is the location of those keys in objective space (e.g. on the table), not the relative position of the keys to me the last time I saw them (e.g. to my left). By contrast, when grabbing for the falling pen, just the opposite features matter: it doesn't matter where the pen is in objective space, it only matters where the pen is with respect to me. Thus, since allocentric and egocentric spatial frames are different, and since translating between them requires time and resources, this too is a conflicting demand between guidance and selection.

Finally, the particular contents required for selection and guidance often differ. In deciding whether to reach for a mug, for example, particular facts about the identity of that mug (e.g. that it's my 'I should be writing' mug) are crucially important. By contrast, when actually reaching for the mug, what matters are precise distances, joint angles, etc. Since computational resources are not infinite, and since computing over more contents requires greater resources and time, these different contents also place conflicting demands on the computational systems that determine them.

Since the demands of the two tasks conflict, a single resource-limited system would be forced to make tradeoffs; favoring either the speed, storage, spatial format, or contents of one other other task, at the cost of the other. A divided system, on the other hand, can allocate separate processing resources to each task, optimizing overall performance. Thus, not only are the local functional claims of Milner and Goodale's model borne out by the evidence (they argue), this functional division itself makes computational sense. And the latter global functional claim entails positive predictions: we should expect the dorsal stream—which is responsible for program guidance—to process stimuli rapidly, to store representations for a relatively short duration, to encode information in an egocentric spatial frame, and to be responsive to contents which are specific to the guidance task. And we should expect the ventral stream—which is responsible for program selection—to process stimuli relatively more slowly, with more detail and tagging, to enable longer-term storage of its representations, to encode information in an allocentric spatial frame, and to process contents which are relative to program selection. Thirty years of neuroscientific research largely confirms these various predictions (Milner & Goodale, 2006).

3.2.3. Two Empirical Challenges to Task-Based Division

I'll now articulate two empirical challenges for task-based division as an account of the function of functional division, multi-stream guidance and multi-stream selection. In section 3.2.1, I rehearsed arguments from previous chapters that evidence of non-proprietary functions is consistent with Milner and Goodale's local functional claims. In this section, I must therefore establish why this same evidence nonetheless challenges their account of the global function of functional division. As we'll see, both multi-stream guidance and multi-stream selection raise the same basic issue for the task-based account: if multiple visual streams with different computational profiles engage in the same basic task (guidance or selection, respectively), then the function of that division cannot be as simple as a difference in the computational demands of those tasks. That is, the empirical evidence now suggests that each of these tasks is carried out using a complex of computational approaches. And a *global* functional account should explain why this is so. Task-based division is structurally incapable of providing such an explanation.

Multi-stream Guidance

The first challenge is multi-stream guidance. While early legion studies suggested that the dorsal stream was uniquely involved in online motor control, increasing evidence suggests that the ventral stream is also involved in the guidance of tasks with sufficient complexity.

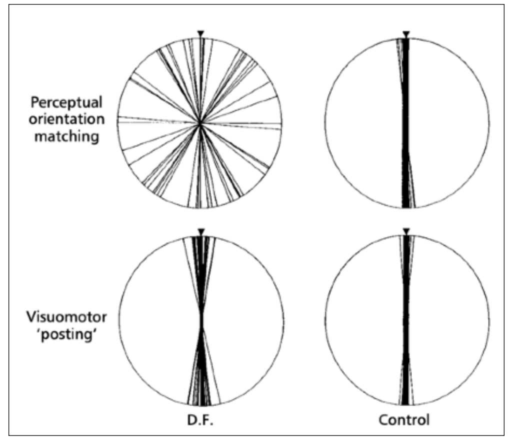

Figure 3.1: Polar plots reflecting the orientation of the cards for D.F. and a control subject in the perceptual matching and visuomotor posting tasks. From Milner and Goodale (2006), data from Goodale *et al.* (1991).

Consider the experiment mentioned in Ch. 2 in which D.F. was asked to both perceptually match the angle of a slot and to post a disc through that slot. In the simplest version in which the slot was a straight line, the results matched the predictions of Milner and Goodale's model: D.F. was unable to accurately match the angle of that slot, but surprisingly adept at posting the disc through that slot. See Fig. 3.1.

In a subsequent experiment, however, Goodale *et al.* (1994) asked D.F. to fit a T-shaped disc through a T-shaped slot, a more complicated task (Fig. 3.2). They found that D.F. could perform this task on only about half of the trials, and she was generally off by about 90 degrees on unsuccessful trials (Fig 3.3). These findings, combined with the above, suggest that dorsal stream

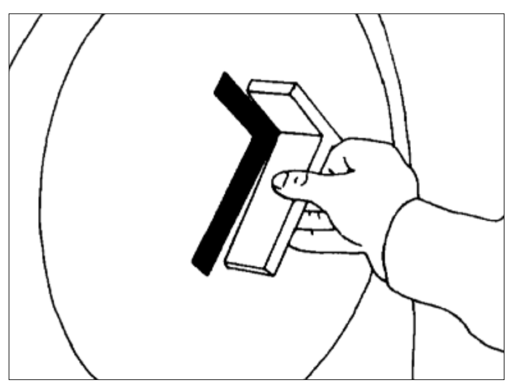

Figure 3.2: From Goodale *et al.* (1994), a representation of the T-shaped posting task.

representations may drive motor guidance in relatively simple aspects of certain tasks (such as navigating the hand around edges), but that ventral stream representations may be involved in the guidance of more complex tasks (such as combining edges into complex object representations). For further such cases, see reviews from McIntosh and Lashley (2008) and Schenk and McIntosh (2010). Drawing on such evidence, Briscoe and Schwenkler (2015) have suggested a kind of dual-systems account of the role of the ventral and dorsal streams in perceptual motor control: the dorsal stream is responsible for the direct control of easy, familiar aspects of tasks, but heavily recruits ventral stream representations in executing novel or difficult motor commands.

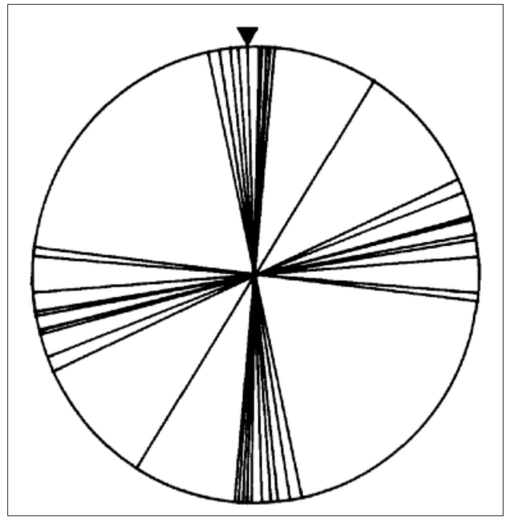

Figure 3.3: From Goodale *et al.* (1994), polar plots reflecting the final orientation in the T-shaped posting task.

Such evidence problematizes task-based division. If there are multiple streams — which by hypothesis have different computational profiles — involved in program guidance, then the separation between the ventral and dorsal streams cannot be justified simply by a difference in the computational requirements of guidance and selection respectively. That is, while task-based division might explain the primary role of the dorsal stream in program guidance, it leaves unexplained the occasional role of the ventral stream in that guidance. What we should want is an account of the function of functional division that helps us both understand why the ventral stream is involved in motor guidance and predict when this should be so. Task-based division is structurally incapable of doing this.

Multi-stream Selection

The problem of multi-stream selection is structurally identical to the problem of multi-stream guidance: if there are multiple streams with different computational profiles performing the role of program selection, then the separation between the streams cannot be justified simply by the conflicting computational demands of selection and guidance.

On the received model, central cognition selects action programs on the basis of (largely conscious) perceptual representations. Thus, multi-stream selection holds if there are multiple streams involved in the formation of those perceptual representations. As mentioned in 2.1, there is now evidence for dorsal stream influence both on perception's object representations and conscious attentional control (Schenk & McIntosh, 2010). And philosophers have further proposed the dorsal stream influences perception's spatial format (Wu, 2014) and 'feeling of presence' (Matthen, 2005). If such influences exist, then it's plausible that dorsal stream processes are at least partially responsible for program selection. Thus, since the ventral stream is also implicated in the formation of the perceptual representations that guide program selection, such influences suggest multi-stream selection.

However, this case of multi-stream selection has two crucial limitations. First, while there is direct evidence for dorsal stream influence on ventral stream representations, substantive influences on either perceptual experience or program selection are largely speculative. That is, studies are not yet refined enough to establish a clear connection between dorsal stream processes and distinctively *perceptual* performance. Second, even if we grant a substantive influence of dorsal stream processes on perceptual representations—and not just pre-perceptual ones—any suggested influence on program selection is necessarily *via* those perceptual representations. That is, the suggested model is one on which both the ventral and dorsal streams contribute to a *unified* perceptual representation which is employed in the selection of action programs. A stronger case for multi-stream selection would posit *independent* influence.

In looking for a program selection stream that challenges task-based division, then, we are looking for four properties. First, of course, the stream must be separate from the ventral stream. Second, it must uncontroversially engage in program selection. That means that the stream must be operative in choosing the broad course of a token behavior, rather than the mere operational details of that behavior. Third, it should engage in program selection independently of ventral stream processes. That is, the strongest case against task-based division will be made if there are multiple streams which determine action programs, not merely a separate stream which influences the ventral stream and downstream cognition in that determination. Fourth, to best undermine the claim of computational division between tasks, it would be best if the stream had a very different set of computational features from the ventral stream. Thus, where the latter is slow, deliberate, and detailed, the target stream should be quick, automatic, and coarse grained.

We find each of these features in the 'low-road' of the subcortical visual system. First, the low-road—which extends from the thalamus to the amygdala, bypassing the visual cortex altogether—is separate from the ventral stream (Tamietto & De Gelder, 2010). Second, since the low-road is responsible for triggering rapid 'fight or flight' responses to threatening stimuli, it is responsible for determining the broad course of behavior, rather than mere operational details (Hamm *et al.*, 2003). Third, the low-road performs this role directly and independently of ventral stream involvement. This independence can be seen, for example, in the speed of low-road processes, which execute well before ventral stream involvement (Tamietto *et al*., 2009). Finally, as this speed suggests, low-road processes have very different computational features from ventral stream ones. Low-road representations are much coarser-grained than ventral stream ones, allowing for rapid and automatic responses to broad patterns in the visual scene (Vuilleumier, Armony, Driver, & Dolan, 2003). Such processing is in sharp contrast to the detailed, deliberate, and slow execution of ventral stream selection.

Thus, while the dorsal stream may be involved in the selection of action programs, the low-road of visual processing presents a stronger challenge to task-based division from multi-stream selection. If the low-road of visual processing independently selects action programs, despite its very different computational profile from the ventral stream, then the separation between streams cannot be justified by a difference in the computational requirements of guidance and selection respectively. Indeed, the evidence above suggests that the computational features of the low-road are more akin to the dorsal stream than the ventral stream. Thus, while task-based division might explain the primary role of the ventral stream in program selection, it leaves unexplained the role of the low-road in selection. What we should want is an account of the function of functional division that helps us both understand why the low road is involved in program selection and predict when this should be so. Task-based division is structurally incapable of doing this.

3.3 Computational Division

Briscoe and Schwenkler (2015) draw on evidence for multi-stream guidance to argue for a new view about the relationship between the ventral and dorsal streams in the guidance of online motor activity. On their view, dorsal stream processes are, as Milner and Goodale contend, primarily responsible for the relatively automatic control of familiar or easy aspects of action tasks, but, contra Milner and Goodale, the ventral stream is heavily recruited in the guidance of tasks which are unfamiliar or difficult. Briscoe and Schwenkler appear to hold, as I've suggested above, that this conclusion is broadly consistent with Milner and Goodale's *local* functional characterization of the dorsal stream — as it maintains that the dorsal stream is organized around program guidance — but inconsistent with their *global* functional characterization of the relationship between the streams — as it precludes a sharp division in streams between program selection and guidance.

Briscoe and Schwenkler's aim is to defend the 'Control Thesis' that "spatial representational contents of visual experience are sometimes used to control actions directed at objects in the surrounding environment" (2015, 1436). Thus, their primary interest is in the *local* function of ventral stream processes in online motor control. And they stop short of offering an explicit account of the *global* function of the division between the ventral and dorsal streams. But their arguments are suggestive of an account, which I'll call "computational division," the view that the functional division between the ventral and dorsal streams is primarily one of computational profile. That is, the dorsal stream (and perhaps the low-road) is engaged in the relatively automatic control of the familiar and easy aspects of visuomotor tasks while the ventral stream is engaged in the more deliberate control of unfamiliar or difficult aspects of those tasks. I'll refrain from attributing this view to Briscoe and Schwenkler, since they don't articulate it, but the similarity between the view they seem to hold and the one defended here deserves note. My aim in this section is to show that computational division addresses many of the shortcomings of task-based division.

Computational division denies any central division in streams between downstream person-level tasks. Instead, it holds that the function of functional division between the visual streams is to allow for multiple computational approaches to be enacted in response to the same downstream tasks. That is, the ventral stream and downstream cognition enact a deliberate, slow, and intelligent computational approach while the dorsal stream and the low road enact a more automatic, quick, and dirty approach. The emphasis here is on automaticity versus non-automaticity, with speed and detail as implications. Thus, the claim is that we have one perception/action system (the ventral stream plus cognition) which has a deliberate (non-automatic) computational profile and other perception/action systems (the dorsal stream and the low road) which have automatic computational profiles. The function of functional division is to allow the enactment of both deliberate and automatic computations toward the same person-level tasks.

To evaluate the plausibility of this account, it will help to clarify what we mean by 'automaticity'. As Ben Phillips has pointed out (2021, manuscript), saying that a process is automatic is vague between saying that that process is mandatory, ballistic, autonomous, or some combination thereof. A process is *mandatory* just in case its triggering conditions entails execution of the process. Unfortunately, for example, given my cooking skills, my smoke detector is mandatory, such that a certain concentration of smoke in the air immediately triggers a building-wide alarm. I'll say that a process is *optional* if it is not mandatory. A process is *ballistic* just in case, once initiated, the process will necessarily run to completion. Also unfortunately, since I often leave out a shirt, my washing machine is ballistic, since it locks upon initiation, preventing alteration until the cycle is complete. I'll say that a process is *interruptible* if it is not ballistic. Finally, a process is *autonomous* just in case, once initiated, its internal operations do not require input from outside. Happily, my water kettle is autonomous, as, once initiated, it automatically brings water to a boil, without requiring further outside input. I'll say that a process is *integrated* if it's not autonomous.

The three notions are each *relative* in the sense that they hold only under certain descriptions of the relevant process. An optional process, for example, is such that multiple independent conditions determine whether the process initiates. But if these inputs are grouped together into broader initiation conditions, then the process is mandatory relative to those conditions. My smart thermostat, for example, will turn on the heater if the room temperature falls below a certain setting, but *only if* my phone is at home. One might say, then, that this thermostat's operation is optional, since it can decide whether to run or not depending on whether I'm at home. Or one could say that its operation is mandatory, since the condition [temperature below setting + Ben's phone is at home] necessarily triggers the heater. Which description is more appropriate depends on one's interests. Moreover, a process is typically mandatory, ballistic, or autonomous only relative to a set of background conditions. My kettle is autonomous, for example, but only if it's plugged in. Finally, it's plausible that each of the three notions come in degrees. My smart thermostat, for example, is more optional than a regular thermostat, but less optional than if I were simply flipping the heater on and off when I wanted to. We might be able to precisify this gradation by appeal to the number of degrees of freedom present in each case. That is, my smart thermostat is more optional than a regular one because more conditions must be fixed to determine whether its operation is triggered.

With these clarifications aside we can now get clearer on the predictions of computational division. On that account, the perception/action streams composed of the ventral stream plus cognition on the one hand and the dorsal stream and the low road on the other differ primarily in their functional profiles. The ventral stream plus cognition is broadly deliberate (optional, interruptible, and integrated) while the dorsal stream and the low road are broadly automatic (mandatory, ballistic, and autonomous). These predictions are borne out by the evidence.

The claim that the ventral stream's influence on action is relatively deliberate is approximately axiomatic.[1] It's widely agreed that the ventral stream's central outputs are to cognition. This explains, for example, the neuroanatomical fact that ventral stream's efferent connections are primarily with prefrontal cortices and the fact that only subjects with intact ventral streams can perform normally on visual judgment tasks (Milner & Goodale, 2006; Goodale & Milner, 2018). And it is a near conceptual truth that cognitive operations are deliberate, in the sense that they are optional, interruptible, and integrated.[2] Cognitive operations are optional, in the sense that, relative to a given stimulus, multiple downstream actions are available. They're interruptible, in the sense that cognition can (at least typically) cancel operations or actions it initiates. And its operations are at least broadly integrated with one another. Since it is broadly agreed that ventral stream influences on action are primarily (or perhaps always) via cognition, it follows that the ventral stream's influence is deliberate.

---

[1] This claim should be distinguished from the claim that ventral stream operations are themselves deliberate, such that, for example, we could simply turn on or off visual face detection. That claim is demonstrably false (e.g. Caldara & Seghier, 2009). The claim here is simply that the ventral stream's *influence on action* is relatively deliberate, in virtue of its cognitive intermediaries.

[2] The claim is that cognition's overall operation is relatively deliberate, in the sense that it is optional, interruptible, and integrated. This is of course consistent with the fact that some cognitive operations are less deliberate and more automatic—as, for example, belief formation may be (Alston, 1988). And it's consistent with the fact that some cognitive processes have relatively automatic processes as their proper parts, or as their inputs. The point is just that cognition is the paradigmatic deliberate system. Exactly how deliberate it is is of course subject to empirical debate.

The case regarding the dorsal stream is more complicated. According to the account outlined above, the dorsal stream's primarily role is the guidance of online motor activity. In reaching for my mug, for example, I employ dorsal stream representations to adjust the aperture of my grip as my hand approaches. That activity, however, is dependent on my already having selected an action program (e.g "grab the mug") which determines the appropriate behavioral response to certain visual parameters. Think of an action program as a set of action commands with conditional structure. The role of the dorsal stream's guidance representation is to indicate which antecedents of those conditions are true and which are false. Thus, overt behavior is determined by combinations of action programs and guidance representations. It follows that dorsal stream representations do not determine overt behavior on their own, but rather only against a backdrop of a particular action program. This complicates our ability to determine the functional profile of dorsal stream representations, since it requires that we assess the automaticity of dorsal stream representations against this backdrop. So the claim is that, relative to a particular action program, dorsal stream guidance representations have their influence on action automatically, rather than deliberately. Specifically, I claim, dorsal stream guidance representations are mandatory, ballistic, and relatively autonomous.

A core source of evidence that dorsal stream processes are automatic is their speed. As described above, dorsal stream processing appears to be organized so as to generate guidance representations quickly: it takes its inputs primarily from fast but coarse magno cells (Livingstone & Hubel, 1988; Merigan & Maunsell, 1993; Sawatari & Callaway, 1996; Milner & Goodale, 2006; Breitmeyer, 2014), its representations can be rapidly updated due to their short storage duration (Milner & Goodale, 2006, section 8.4.3), and so forth. In general, deliberate processes are more time consuming than automatic processes, as the former operate over more inputs, must weigh further options, etc. Thus, rapid processing is a core piece of evidence that a process is automatic. In particular, it's likely that the dorsal stream's influence on action is both highly mandatory and highly ballistic.

The case for the dorsal stream's autonomy is less clear-cut. As rehearsed in Section 3.2, there is substantial evidence of functional and informational connectivity between the ventral and dorsal streams. The latter in particular suggest that dorsal stream processes operate in part over representations in the ventral stream, undermining the isolability of the two streams. Thus, it's not plausible that dorsal stream guidance representations are fully autonomously produced. However, it remains plausible that dorsal stream processes are *more autonomous* than ventral stream ones, as, once again, the latter influence action only via cognition, a paradigmatically integrated mental unit. Thus, the processes which influence action via the ventral stream are likely far more integrated than those that influence action via the dorsal stream.

Our final claim is that the low-road is automatic. Specifically, low-road influence is mandatory, and relatively autonomous and ballistic. The primary evidence for automaticity regarding the low-road is, again, speed. Low-road processes are causally antecedent to ventral stream ones, as they proceed directly from LGN to the amygdala, bypassing visual cortex (Tamietto & De Gelder, 2010). And evidence supports the implication that low-road processes are much faster than other perceptual processes, triggering fight or flight responses prior to conscious awareness (Tamietto *et al*., 2009). This speed weighs in favor of low-road processes being both mandatory and ballistic.

The question of autonomy is independent of these two issues, and turns on whether top-down processing influences perceptual representations in the LGN. While there has been little exploration of such influence, there is some evidence that task-demands cause back-propagation from visual cortex to LGN, enhancing or suppressing stimulus processing (Casagrande, Sáry, Royal, & Ruiz, 2005). That is, one's goals influence the allocation of attentional resources as early as the thalamus. This suggests some indirect influence of 'higher' representations on low-road representations. Thus, low-road processes are not entirely autonomous. Such influences, however, appear to be quite sparse, suggesting far less integration than ventral stream processes. Thus, we can say that low-road perceptual representations appear to be relatively autonomous.

There's also some evidence for diachronic influence of cognition on amygdala activity. Specifically, there is now widespread evidence that activity in the medial prefrontal cortex is anti-correlated with activity in the ventral amygdala (Milad & Quirk, 2002; Milad, Vidal-Gonzalez, & Quirk, 2004; Santini, Ge, Ron, Ortiz, & Quirk, 2004; Holmes *et al.*, 2012). This suggests that top-down influence can downregulate amygdala function, limiting its influence on threat response. Such influence suggests that the low-road is at least partially and diachronically interruptible. We cannot necessarily prevent representations in the LGN from triggering responses in the early amygdala, but we can regulate the downstream influence of that triggering on behavior. Here again, however, such interruptibility is to be contrasted with cognitive function, which is *synchronically* interruptible to a substantial degree. Thus, we can say that low-road activity is *relatively* ballistic.

Evidence across the three streams, then, supports the central contentions of computational division. Because the ventral stream has its influence on action via central cognition, that influence is paradigmatically deliberate: optional, interruptible, and integrated. By contrast, the dorsal stream and the low road influence action directly, and their computational profiles are relatively automatic: broadly mandatory, ballistic, and autonomous.

Moreover, computational division can explain the evidence that challenged task-based division. The latter holds that different person-level tasks—program guidance and selection respectively—place conflicting computational demands on the systems that underlie them, thus requiring multiple streams with different computational profiles dedicated to each task. As we saw, however, the evidence suggests that multiple streams—each with different computational profiles—are involved in each of these person-level tasks. Thus, the function of functional division cannot be as simple as allowing for different systems for different tasks. But computational division holds that there are *task-neutral* advantages to enacting both automatic and deliberate processes in the performance of tasks. Thus, it predicts both multi-stream guidance and multi-stream selection. And this is precisely what we find: the ventral and dorsal streams enact deliberate and automatic processes respectively in the joint determination of program guidance, and the ventral stream and the low road enact deliberate and automatic processes respectively in the joint determination of program selection. This allows for both tasks to be performed using combinations of both quick-and-dirty and slow-and-deliberate strategies, optimizing performance.

3.4 Coupled and Decoupled Processes

Computational division captures the broad division of labor between the visual streams. And it explains the otherwise puzzling facts of multi-stream selection and multi-stream guidance. But, as a final step in the development of the view defended here, I'd like to revisit whether we should think of the function of functional division as task-neutral. Is the fundamental division between the ventral, dorsal, and low road streams simply one between deliberate and automatic processes? I think it is not, and the way in which it is not moves us closer to a middle ground between the two views discussed thus far.

The view that I'm about to defend is motivated by three questions that remain for computational division. First, if the fundamental distinction between the streams is one between automatic and deliberate computational approaches, why are there at least three streams, rather than two? Second, given that the dorsal stream and the low road both enact broadly automatic processes, what is the function of the division between them? That is, why not just have one automatic visual processing system just as there appears to be one broadly deliberate one? Third, and on the other hand, why does there appear to be just one deliberate processing stream, which operates across multiple task-types, when there are multiple automatic processing streams?

My view is that the division between visual streams is primarily one between a *task-decoupled* stream and multiple *task-coupled* ones. Thus, the function of functional division is to enable the deployment of both task-coupled and task-decoupled processes in response to the same kinds of tasks.

A task-coupled process is one such that fulfillment of its triggering conditions entails fulfillment of a downstream person-level task. All coupled processes are automatic, in that, at least, they are both mandatory and ballistic (and typically autonomous), but not all automatic processes are coupled, since not all automatic processes have a stereotyped person-level influence. Face detection, for example, is broadly automatic (it's mandatory, ballistic, and broadly autonomous), but since the detection of a face is not associated with a particular person-level task, it is not coupled.

This notion of 'coupling' is inspired by Kim Sterelny's (2003) account of the evolution of perception/action systems. His idea is that such systems began as simple coupled systems. Cockroaches, for example, have special antennae which detect gusts of wind as they pass over the insect. When predatory toads leap toward the cockroach, the corresponding gust of wind automatically triggers a fleeing response. In this way, a simple coupled system allows the cockroach to avoid predation at minimal computational and metabolic cost (Camhi, 1984; Sterelny, 2003). But Sterelny argues that hostility (predation, competition, etc.) tends to select away from such easily exploitable coupled systems and toward more flexible decoupled ones. Thus, highly evolved critters like us will have highly complex perception/action systems such that token perceptual inputs determine intermediate representations, which are then employed to make highly flexible decisions about what to do next.

Because decoupled systems are such that a given perceptual stimulus can be associated with a wide variety of internal or external behaviors, their representations are necessarily more abstract. We might think of coupled systems as employing representations which are roughly imperative: they simply instruct action systems to do particular things—i.e. to either run an action program or to update the guidance of that program. Decoupled systems, by contrast, operate over representations which are deployable across a range of tasks. Thus, we can think of such systems as employing descriptive models of the distal environment which are employed in making decisions about what to do next. These various features hang together: a decoupled system has options about what it can do in response to a given stimulus; thus, its representations are abstracted away from any given downstream task; and thus, it must employ those representations to make decisions about what to do next. Because these various features hang together in this way, Sterelny argues that highly decoupled systems are the foundation of cognition.

If this line of thought is correct, then claims about the computational nature of perception/action systems cannot be fully abstracted from claims about the range of tasks those systems can engage with. Automatic systems generally specialize to particular task-types, employing only locally useful information to allow for rapid and efficient influence on action — that is, they're coupled. And deliberate systems become deliberate by gradually abstracting away from local task demands to enable greater flexibility in response — that is, they're decoupled.

I claim that the ventral stream and downstream cognition constitute a decoupled perception/action stream, while the dorsal stream and the low road constitute coupled perception/action streams. That claim is a kind of middle ground between task-based division and computational division. With task-based division, I accept a broad division between 'vision for perception' — here recast as visual processing which generates abstract models for the broad guidance of cognition — and 'vision for action' — here understood as action-coupled processing for particular downstream tasks (program selection in the case of the low road and program guidance in the case of the dorsal stream). And with computational division, I accept a division between broadly deliberate processing — carried out by a task-neutral system composed of the ventral stream and cognition — and automatic processing — carried out by multiple task-specific systems including the dorsal stream and the low road.

Like computational division, my account can explain the evidence that challenged task-based division: it directly predicts that there should be both coupled and decoupled processing streams for the same downstream task, when such a division makes computational sense. And this is precisely what we find in the models: it makes sense that there should be a coupled stream dedicated to the familiar, easy aspects of program guidance, while also allowing for a more deliberate stream which assists in the guidance of unfamiliar or difficult aspects of program guidance (Briscoe & Schwenkler, 2015). And it makes sense that there should be a coupled stream dedicated to trigger rapid fight-or-flight action programs in response to threatening stimuli, even as a more intelligent, deliberate stream handles the selection of other programs.

Unlike simple computational division, however, my account answers the three questions with which I began this section. Why are there at least three streams, rather than just two, each corresponding to a particular computational profile? Because coupled processing systems are typically dedicated to particular task-types. Thus, in addition to a general decoupled processing system, it makes sense to have multiple coupled ones. What is the function of the division between the dorsal stream and the low road, given that they have the same general computational profile? The division allows for each stream to specialize in a particular task-type, optimizing performance. And why is there just one deliberate processing stream, when there are multiple automatic processing streams for different tasks? Because decoupled systems are, by their nature, task-neutral. It thus makes sense to employ one general model in making decisions across a range of task-types.

I mentioned in the introduction that local and global functional questions are mutually constraining: global functional accounts must be consistent with facts about the local function of individual mental units, and local functional accounts of those units are informed and revised by plausible answers to global functional questions. We've seen this dynamic play out in this chapter, and I'd like to conclude by highlighting several instances.

First, of course, accounts of global function are beholden to local functional facts. As I've just argued, two otherwise plausible accounts of the function of functional division are undermined by the empirical evidence. Because task-based division holds that the division between visual streams enables each stream to specialize in a particular task according to its computational requirements, it is undermined by evidence for multi-stream guidance and selection. And because task-neutral computational division holds that the central division between streams is one of computational profile, not downstream task, it is undermined by evidence that the dorsal stream and the low road are task-specialized. In their place, I argued for a kind of middle view which predicts multiple task-coupled processing streams (the dorsal stream and the low road) and a single task-decoupled processing stream (composed of the ventral stream and downstream cognition).

Second, our discussion highlights the importance of differentiating between global and local functional claims. Conflating the two can cause us to inappropriately reject accurate local functional ascriptions. Specifically, as we saw in Section 3.2, Milner and Goodale's global functional claim — that the function of functional division is simply to enable task-specialization — but not their local functional claims — that the dorsal stream is for motor guidance while the ventral stream is for conscious perception — is undermined by evidence for multi-stream guidance and selection. In fact, my own account of the function of functional division is largely consistent with the received accounts of the local function of each stream. Thus, the evidence which is typically taken to undermine Milner and Goodale's theory in fact only undermines their account of the function of functional division. And once we replace their account with the correct one, we see that the central tenets of their model, along with their core implications, remain.

Finally, however, answers to global functional questions can help us refine local functional claims, and to discover new features of the visual streams. Two implications of the account defended here are worth emphasizing here. First, my account supports the contention that the low road constitutes a third, independent visual stream. That is, independent of the direct case for the low-road as visual stream (which I've only gestured at here), my account of the function of functional division predicts that, *ceteris paribus*, there should be a coupled processing stream whenever the computational demands of a person-level task support it. The need for rapid program selection in response to threats is an obvious case of such a demand. Thus, the account provides broader theoretical support for the existence of such a stream.

Second, my account makes predictions about the ways in which we can influence operations in the respective streams. Because the ventral stream's influence is via cognition, our influence on downstream action is *direct.* This is a natural implication of decoupled processing streams: because there is no immediate connection between perceptual input and downstream behavior, there's room for influence on that output. The dorsal stream and the low road, by contrast, are coupled systems; thus, their impact on behavior is automatic and thus less susceptible to direct, top-down control. This suggests that our influence over their behavior is at best indirect, involving broader regulation of ourselves and our environment. This final point reiterates the most important implication of the division in visual streams: while our pre-theoretic intuitions give a central role to cognition in the determination of visually-guided behaviors, top-down control is in fact both incomplete and disunified. Thus, we must enact a variety of strategies to influence such behaviors.

References

Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical perspectives*, *2*, 257-299.

Breitmeyer, B. G. (2014). Contributions of magno- and parvocellular channels to conscious and non-conscious vision. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1641), 20130213.

Briscoe, R., & Schwenkler, J. (2015). Conscious Vision in Action. *Cognitive Science*, *39*(7), 1435–1467.

Caldara, R., & Seghier, M. L. (2009). The fusiform face area responds automatically to statistical regularities optimal for face categorization. *Human brain mapping*, *30*(5), 1615-1625.

Camhi, J. (1984). *Neuroethology*. Sunderland, MA.

Casagrande, V. A., Sáry, G., Royal, D., & Ruiz, O. (2005). On the impact of attention and motor planning on the lateral geniculate nucleus. *Progress in brain research*, *149*, 11-29.

Cloutman, L.L.. (2013). Interaction between dorsal and ventral processing streams: where, when and how? *Brain Lang*. 127, 251–263.

Goodale, M. A., Jakobson, L. S., Milner, A. D., Perrett, D. I., Benson, P. J., & Hietanen, J. K. (1994). The nature and limits of orientation and pattern processing supporting visuomotor control in a visual form agnosic. *Journal of Cognitive Neuroscience*, *6*, 46–56.

Goodale, M. and Milner, A. (2004). *Sight Unseen: An Exploration of Conscious and Unconscious Vision.* Oxford University Press.

Goodale, M. A., & Milner, A. D. (2018). Two visual pathways – Where have they taken us and where will they lead in future? *Cortex*, *98*, 283–292. https://doi.org/10.1016/j.cortex.2017.12.002

Grafton, S.T., (2010). The cognitive neuroscience of prehension: recent developments. *Exp. Brain Res*. 204, 475–491.

Haak, K. V., & Beckmann, C. F. (2018). Objective analysis of the topological organization of the human cortical visual connectome suggests three visual pathways. *Cortex*, *98*, 73-83.

Hamm, A. O., Weike, A. I., Schupp, H. T., Treig, T., Dressel, A., & Kessler, C. (2003). Affective blindsight: intact fear conditioning to a visual cue in a cortically blind patient. *Brain*, *126*(2), 267-275.

Holmes, A., Fitzgerald, P. J., MacPherson, K. P., DeBrouse, L., Colacicco, G., Flynn, S. M., ... & Camp, M. (2012). Chronic alcohol remodels prefrontal neurons and disrupts NMDAR-mediated fear extinction encoding. *Nature neuroscience*, *15*(10), 1359-1361.

Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, 240, 740–9.

Matthen, M. (2005). *Seeing, doing, and knowing: A philosophical theory of sense perception.* Oxford University Press.

McIntosh, R. D., & Lashley, G. (2008). Matching boxes: Familiar size influences action programming. *Neuropsychologia*, 46, 2441–2444.

Merigan, W. H., & Maunsell, J. H. (1993). How parallel are the primate visual pathways?. *Annual review of neuroscience*, *16*(1), 369-402.

Michel, F., & Henaff, M. A. (2004). Seeing without the occipito-parietal cortex: Simultagnosia as a shrinkage of the attentional visual field. *Behavioural Neurology*, 15, 3–13.

Milad, M. R., & Quirk, G. J. (2002). Neurons in medial prefrontal cortex signal memory for fear extinction. *Nature*, *420*(6911), 70-74.

Milad, M. R., Vidal-Gonzalez, I., & Quirk, G. J. (2004). Electrical stimulation of medial prefrontal cortex reduces conditioned fear in a temporally specific manner. *Behavioral neuroscience*, *118*(2), 389.

Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action* (2nd ed). Oxford University Press.

Mishkin, M., & Ungerleider, L. G. (1982). Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural Brain Research*, *6*(1), 57–77. https://doi.org/10.1016/0166-4328(82)90081-X

Perenin, M. T., & Vighetto, A. (1988). Optic ataxia: A specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain*, *111*, 643–674.

Phillips, B. (2021). Seeing Seeing. *Philosophy and Phenomenological Research*, *102* (1), 24-43.

Phillips, B. (manuscript). Automaticity and Social Perception.

Pisella, L., Sergio, L., Blangero, A., Torchin, H., Vighetto, A., & Rossetti, Y. (2009). Optic ataxia and the function of the dorsal stream: contributions to perception and action. *Neuropsychologia*, *47*(14), 3033-3044.

Santini, E., Ge, H., Ren, K., de Ortiz, S. P., & Quirk, G. J. (2004). Consolidation of fear extinction requires protein synthesis in the medial prefrontal cortex. *Journal of Neuroscience*, *24*(25), 5704-5710.

Sawatari, A., & Callaway, E. M. (1996). Convergence of magno-and parvocellular pathways in layer 4B of macaque primary visual cortex. *Nature*, *380*(6573), 442-446.

Schenk, T., & McIntosh, R. D. (2010). Do we have independent visual streams for perception and action? *Cognitive Neuroscience*, *1*(1), 52–62.

Sterelny, K. (2003). *Thought in a Hostile world*. Blackwell Publishing.

Striemer, C., Blangero, A., Rossetti, Y., Boisson, D., Rode, G., Vighetto, A., et al. (2007). Deficits in peripheral visual attention in patients with optic ataxia. *Neuroreport*, *18*, 1171–1175.

Striemer, C. L., Chapman, C. S., & Goodale, M. A. (2009). "Real-time" obstacle avoidance in the absence of primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 15996–16001.

Striemer, C., Locklin, J., Blangero, A., Rossetti, Y., Pisella, L., & Danckert, J. (2009). Attention for action? Examining the link between attention and visuomotor control deficits in a patient with optic ataxia. *Neuropsychologia*, *47*, 1491–1499.

Tamietto, M., Castelli, L., Vighetti, S., Perozzo, P., Geminiani, G., Weiskrantz, L., & de Gelder, B. (2009). Unseen facial and bodily expressions trigger fast emotional reactions. *Proceedings of the National Academy of Sciences*, *106*(42), 17661-17666.

Tamietto, M., & De Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature Reviews Neuroscience*, *11*(10), 697-709.

van Honk, J., Tuiten, A., van den Hout, M., Koppeschaar, H., Thijssen, J., de Haan, E., & Verbaten, R. (1998). Baseline salivary cortisol levels and preconscious selective attention for threat: A pilot study. *Psychoneuroendocrinology*, *23*(7), 741-747.

van Polanen, V., & Davare, M. (2015). Interactions between dorsal and ventral streams for controlling skilled grasp. *Neuropsychologia*, *79*, 186–191.

Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature neuroscience*, *6*(6), 624-631.

Wu, W. (2014). Against Division: Consciousness, Information and the Visual Streams. *Mind & Language*, *29*(4), 383–406. https://doi.org/10.1111/mila.12056

Chapter 4: Responsibility and Perception

On June 2nd, 2010, Armando Galarraga nearly became the twenty-first pitcher in Major League Baseball history to pitch a perfect game—one in which no opposing batter reaches base. The feat is rare; only twenty-three of the approximately 217,000 MLB games have been perfect games. On what would have been Galarraga's final pitch, Jason Donald hit a ground ball that first baseman Miguel Cabrera ran to retrieve. Galarraga raced to first base to catch Cabrera's throw and force Donald out. Umpire Jim Joyce ruled Donald safe, ending the perfect game. But that ruling was incorrect. Video clearly shows that Galarraga caught the ball and touched the base before Donald. After the game, Joyce immediately apologized.

All agree that Joyce's ruling was incorrect. But those familiar with sports fans will recognize additional language that is used in such cases: Joyce was 'blind', an 'idiot', and so on. The implication is that Joyce was not only wrong, but irrational, irresponsible, or unjustified in ruling as he did. That assessment is not without warrant. In what's called a 'force play', Galarraga tried to catch the ball and touch the base before Donald. When close, it's difficult to tell who touched the base first or whether the fielding player caught the ball in time. But this case was not close; Galarraga caught the ball and touched the base while Donald was a full stride away. This should have been an easy call for an umpire with Joyce's experience. Umpires are taught a rule for perceiving force plays: they should look at the base while listening for the sound of the ball hitting the glove. This makes far easier the otherwise difficult task of tracking the locations of the ball and two people's feet. One hypothesis is that Joyce failed to follow this rule.

While fans often jeer at umpires that make incorrect rulings, criticism has particular force when umpires miss something that is clearly evident. That is, criticism of Joyce's ruling has a distinctively epistemic element. If this criticism is justified by Joyce's failure to follow a rule about how to best perceive force plays, then this case adds to growing evidence for the epistemic relevance of an experience's etiology (Siegel 2013). The case also suggests a new dimension of epistemic evaluation of that etiology. Joyce is accused of failing to follow a rule—one that he should have followed—about how to position himself mentally and physically in order to have the right kind of perceptual experience. That is, Joyce formed his experience irresponsibly; and this redounds on the epistemic standing of his resulting belief. This consideration is independent of more familiar epistemic influences on perceptual beliefs—such as reliability or inferential character.

In this chapter, I articulate and defend this *sui generis* dimension of epistemic influence on perceptual belief. I argue that beliefs based on irresponsibly formed experiences—experiences whose causes were not appropriately regulated by the subject—are doxastically unjustified. In Section 4.1, I articulate the view I'm defending. In Section 4.2, I defend the premise that Joyce's belief, but not that of a similarly situated novice, is unjustified. In Section 4.3, I show that this difference is best explained by irresponsible experience formation. In Section 4.4, I show that the epistemic relevance of responsible experience formation has broad implications for the epistemology of perceptual beliefs.

4.1 Responsible Experience Formation

Joyce failed to follow a rule—one that he should have followed—about how to best perceive force plays. I'll argue that Joyce formed his experience irresponsibly and that this redounds on the justification conferred to his belief. Specifically, I'll defend two principles:

> Responsibility: S is responsible in having an experience regarding *p* just in case S does what's reasonably expected to ensure that S's experience is most apt regarding *p*.
>
> Downgrade: If S forms a first-order belief B with content P on the basis of an experience which is irresponsibly formed with respect to P, B is thereby doxastically unjustified, assuming that S has no other basis for B.

Responsibility appeals to 'reasonable expectations' on experience formation. I'll assume that one is reasonably expected to ϕ (in this sense) just in case one has an epistemic obligation to ϕ. My view thus extends such obligations to the formation of experiences. My argument will rely on there being both general and special epistemic obligations—and thus reasonable expectations—regarding the formation of experiences. A general obligation to attend to immediately available relevant evidence is an example of the former. Crucially, I'll claim that Joyce's obligation to follow the rule is a special obligation required of Joyce in virtue of his being a professional umpire. More on this below.

Responsibility also appeals to an experience's 'aptness'. Experiences have multiple good-making features regarding a particular content. They are, for example, more or less informative regarding *p* and more or less reliable in these verdicts. The most apt experience regarding *p* is the one which best trades off these different virtues. Aptness doesn't reduce to veridicality, since different veridical perceptions can be more or less informative or reliable regarding *p*. Nor does responsibility reduce to aptness, since Responsibility concerns what one *does* to ensure aptness. One can have the most apt experience—via luck—without forming that experience responsibly. And one can form an experience responsibly—but unluckily—and fail to form the most apt experience.

Responsibility is based on Ru Ye's (2018) formulation of responsible belief formation. On her view, forming a belief irresponsibly—such as by failing to appropriately respond to a defeater—undermines doxastic justification. I extend this idea to the formation of experiences. I've based my formulation on Ye's because, unlike other recent accounts such as Peels's (2017), her formulation is largely neutral about which particular epistemic obligations we have, the conditions under which we have those obligations, and the connection between those obligations and 'normative attitudes' like praise or blame. I'll argue that in conditions like Joyce's, experts have a special epistemic obligation to enact certain procedures in forming their experiences. And I'll claim that in such cases we can reasonably blame experts who fail to meet such obligations. But rather than offering a general theory of the particular obligations we have and the conditions under which we have them, my aim is to show that my view has unique explanatory advantages in cases like Joyce's. I'll thus remain neutral regarding a broader theory (though see Section 4.4).

Downgrade is based on Susanna Siegel's (2013) 'Doxastic Downgrade Thesis' (more on which in Section 4.3.3). The restriction to first-order beliefs excludes beliefs like "It looks as if there is a cup in front of me", which describe the experience itself and are not intuitively subject to downgrade. While Siegel has gone on to claim that downgrade impacts an experience's ‘epistemic charge’ such that both an experience and its etiology are rationally assessable (2017), Downgrade is neutral on these points. It allows, for example, that downgrade occurs because the *subject*—and not the etiology or experience—is epistemically evaluable.

Responsibility and Downgrade are relativized to both particular subjects of experience and particular belief contents. These restrictions allow for an experience to confer justification for some beliefs (such as, “Donald ran to first base”) without conferring justification to others (such as, “Donald touched the base immediately before Galarraga”). One potential worry for my view is that in a less-close case—such as when Donald reaches the base well before Galarraga—Joyce can form his experience irresponsibly—by failing to follow the rule—but remain justified in his belief that Donald was safe. But this worry is addressed by the restriction to particular contents: failure to follow the rule only undermines the justificatory force of experiences regarding the close cases to which it applies, not all experiences of force plays. However, the restriction is neutral on whether experiences have contents (which is controversial; see Section 4.3.3). It requires only the uncontroversial claim that the same perceptual experience can convey different levels of justification for different belief contents.

My view, then, is that, since Joyce is a professional umpire, he has a special epistemic obligation—and thus is reasonably expected—to follow the rule when perceiving force plays. Since he formed his experience about a force play without following the rule, it follows from Responsibility that he formed that experience irresponsibly. And since he formed his belief that Donald was safe on the basis of the irresponsibly formed aspect of his experience, it follows from Downgrade that that belief is unjustified. In the next two sections, I defend this analysis by showing that it best explains our verdicts in Joyce-like cases. Before turning to the direct argument for this claim, however, I must clarify the explanandum. I turn to this now.

4.2 Perceptual Expertise and Special Epistemic Obligations

In the next section, I'll argue that irresponsible experience formation, and not defeat, reliability, or inference, can explain why Joyce's belief, but not that of a similarly situated novice, is unjustified. In this section, I defend this difference in justification by arguing that it alone can explain why we can reasonably blame Joyce, but not the novice, for forming the belief that Donald is safe. I claim that Joyce, but not the novice, has a special epistemic obligation to follow the rule, and that his failure to do so undermines the justificatory status of his resulting belief.

It's worth noting, however, that one could reasonably (but wrongly, I think) deny the existence of such special obligations while accepting both Responsibility and Downgrade. Doing so would result in the view that Joyce and the novice either both responsibly or both irresponsibly formed their experiences and thus are equally justified in their resulting belief. This view lacks the corresponding explanatory advantages—since, as we will see, the verdict that Joyce and the novice are equally justified *can* be explained by defeat, reliability, or inference—but it could be defended by appeal to a more general theory of justification. Theories which explain doxastic justification by appeal to responsible *belief* formation—including deontological (Kornblith, 1983; Peels, 2017; and Ye, 2018) and virtue epistemic (Code, 1987; Montmarquet, 1993; Zagzebski, 1996; Baehr, 2011) approaches—are natural allies to the view defended here. Defenders of such theories may well subscribe to Responsibility and Downgrade without subscribing to special epistemic obligations of the sort posited here. Indeed, as I'll discuss in Section 4.4, Responsibility and Downgrade have theoretical interest well beyond their ability to explain special epistemic obligations regarding the formation of experiences.

But a direct defense of Responsibility and Downgrade from the existence of special epistemic obligations—rather than an indirect argument from a general theory of justification—has three advantages. First, since a direct defense is neutral between general theories of justification, it can be accepted independently of such theories. Second, since a direct defense can be accepted independently of general theories, it can serve as an independent constraint on them. If the argument of this chapter is sound, then any general theory of doxastic justification must explain the epistemic import of irresponsible experience formation just as, for example, they must explain the epistemic import of defeaters. I'll suggest that the view defended here is incompatible with phenomenal conservatism (a paradigm internalist view) and compatible only with modified versions of process reliabilism (a paradigm externalist view). Finally, while previous defenses of responsibilist theories have focused on capturing the normative force of existing influences on doxastic justification—such as higher-order defeaters (Ye 2018)—the direct defense offered here cements responsible experience formation as a *sui generis* kind of influence, directly akin to defeat or bad inference. Responsible experience formation is an independent epistemic phenomenon that deserves analysis from within multiple competing frameworks.

I turn now to a defense of the explanandum. Suppose that a novice also attempts to determine whether Galarraga's force play was successful. Unlike Joyce, the novice doesn't know the rule about perceiving force plays. In the same circumstances as Joyce's, they thus break the rule and also come to falsely believe that Donald was safe. This chapter aims to explain why the novice's belief, but not Joyce's, is justified.

Since Joyce is a professional umpire, I'll say that he is a *perceptual expert* regarding force plays. A perceptual expert about a domain is someone that we reasonably expect to form better experiences—e.g. to attend to the right things—in that domain. Examples abound. Doctors are perceptual experts regarding x-rays. Police officers are perceptual experts regarding traffic violations. Chicken sexers are perceptual experts regarding chicken sex.

There are at least two conditions that might make someone a perceptual expert. First, when someone has proprietary knowledge about perceiving in a domain, we reasonably expect them to follow it. Since Joyce knows how to correctly perceive force plays, it's reasonable to expect that he will. Second, when someone occupies a special epistemic social role, it's reasonable to expect them to follow standard rules for occupiers of that role (Feldman 2000, 676). Since Joyce is a professional umpire, it's reasonable to expect him to follow standard perceptual rules for umpires. Sometimes these two conditions might come apart, as when an incompetent umpire fails to know the rule or when an enthusiast novice does. And we might wonder what we can reasonably expect in such cases. But since the conditions go together in the Joyce case, and since my argument requires only one positive case, I won't weigh in on this issue here.

I claim that being a perceptual expert creates special obligations such that having expertise in a domain raises the standards of justification in that domain. Since both fail to meet this higher standard, Joyce's belief, but not the novice's, is unjustified. I posit this difference because it uniquely explains why we can reasonably blame Joyce, but not the novice, for acting as he does. I'll consider two alternative explanations.

A first alternative explanation for why we can blame Joyce (and not the novice) is that he fails a *practical*, rather than epistemic, obligation. Such a failure need not redound on doxastic justification. On this view, both Joyce and the novice are justified in their belief that Donald was safe, but Joyce has a special obligation to *act* on such justified beliefs only when he's followed the rule. The recent epistemic literature has proposed several cases which fit this general profile. Consider:

> Jonathan is a quality inspector at the widget factory. One in a million widgets produced at the factory is defective. Jonathan's job is to randomly inspect widgets to confirm that they work. Before inspecting a widget, he can justifiably believe that that widget is working, since only one in a million is defective. But before inspecting the widget, it would be inappropriate for Jonathan to write down 'working' on his form. And we could reasonably blame him if he did so. (Simplified from Adler, 2004)
>
> Matilda is Derek's oncologist. She receives reliable testimony from a colleague—who has reviewed Derek's tests, but who doesn't share the details—that Derek has pancreatic cancer. Thus, Matilda is justified in believing that Derek has pancreatic cancer. But it would be inappropriate for Matilda to *assert*, based on this evidence, that Derek has cancer, since assertion in this context requires that Matilda have more information about the facts of the case. (Simplified from Lackey, 2011)

In both cases, while the subject's belief is justified, they can be blamed if they don't enact additional procedures before acting on their belief. Could we give a similar explanation of the Joyce case?

There are two crucial differences between such cases and the Joyce case. First, the arguments in which the above cases are embedded rely heavily on the fact that the subject acts on a belief with the wrong *kind* of justification. Jonathan shouldn't rely on statistical evidence when determining whether a particular widget is faulty. Matilda shouldn't rely on isolated testimonial justification when asserting that Derek has cancer. It's now commonly acknowledged that certain kinds of justification—especially 'indirect' forms like statistical or testimonial justification—can be insufficient for epistemically appropriate action. Joyce, by contrast, should (and must) rely on perceptual information when determining whether Donald is safe. He is blameworthy because of the *way* in which he forms his experience, not the fact that he bases his ruling on it at all. There's no mismatch between this *type* of justification and Joyce's action. Without such a mismatch, a direct impact on doxastic justification is the more plausible explanation.

Second, while most will agree that Jonathan and Matilda have the appropriate doxastic attitude given their circumstances, Joyce has the wrong doxastic attitude given his. Regardless of his subsequent ruling, given that he failed to follow the rule, he should suspend judgement about whether Donald was safe, just as a scientist who fails to enact strict experimental controls should suspend judgement about the generalizability of their results. If Joyce should suspend judgement, then his belief is unjustified.

A second explanation is that there is a *general* epistemic obligation to follow the rule in perceiving forced plays, but the novice is excused from blame for failing to follow it. On this view, both Joyce and the novice have unjustified beliefs, but only Joyce is blameworthy.

Notice that, for such a response to work, it must deny the existence of *any* special obligations of the sort posited here. To refute this chapter's argument, it's not enough to simply deny a difference in justification in the Joyce case; one must also deny that there are any cases *like* the Joyce case on which there is such a difference. But there are cases like the Joyce cases for any instance of expertise. Doctors are taught *how to look* at x-ray results; police officers are taught *how to look* for traffic violations; chicken sexers are taught *how to look* for male and female chicks. Beyond perceptual cases, scientists are taught how to conduct proper experiments; judges are taught how to attend to the details of the law; etc. For the present suggestion to work, there must either be a general epistemic obligation to enact the relevant procedure in these disparate domains, or such know-how is epistemically irrelevant (except as a potential source of defeaters). Either result seems implausible. On the first horn, it seems implausible, for example, that we are born with a general epistemic obligation to perceive in accordance with the rule regarding force plays. On the second horn, it seems implausible that learning about such procedures has no impact on *prima facie* justification. A more plausible picture—the one defended here—is that epistemic progress—such as discovering better procedures for perceiving in specialized domains—brings with it special epistemic obligations for experts to enact those procedures. Those obligations directly impact doxastic justification.

4.3 Competing Explanations

I've just argued that the novice's belief, but not Joyce's, is justified. The view presented in Section 4.1 gives a straightforward explanation of this difference: since Joyce, but not the novice, is reasonably expected to follow the rule, it follows from Responsibility that the novice, but not Joyce, formed their experience responsibly. And since they both formed their belief that Donald was safe on the exclusive basis of this experience, it follows from Downgrade that the novice, but not Joyce, is justified. This explanation is non-traditional in a variety of ways. Where accounts of empirical justification typically begin with perceptual experiences (and other kinds of mental states), the view defended here makes room for the direct epistemic relevance of the formation of perceptual experiences; including how extra-bodily processes contribute to those experiences. While some have defended the epistemic relevance of the formation of experiences (e.g. Siegel 2017), they have typically done so by appeal to top-down effects from other mental states, not the direct control of the formation of experiences themselves. In short, responsible experience formation is a different *kind* of influence on the doxastic justification of perceptual beliefs. Thus, if the Joyce cases are to motivate the positing of such an influence, we must first rule out a more traditional explanation of the difference in justification between Joyce and the novice. I'll consider three alternative explanations—from defeat, reliability, and inference—and show they fail. The upshot is that responsible experience formation is an independent—though not necessarily exclusive—determinant of the epistemic status of perceptual beliefs.

4.3.1 Defeat

A common position—particularly among internalists—is that the etiology of experience is not directly relevant to justification, but that evidence *about* that etiology can undermine otherwise justified perceptual beliefs (e.g., Pollock & Cruz 1999). If, for example, I seem to see a red cup sitting on a table, then I am *prima facie* justified in believing that the cup is red. But if I learn that the cup is illuminated by red lights, then I cannot rule out its being a white cup that merely looks red in these conditions. Thus, my belief is defeated by my knowledge of the illumination. It has recently been proposed that a fact is a 'normative defeater' if, (a) it would function as a classical defeater if known and (b) the subject *should* know it (Lackey, 2008, 45). If, for example, I ignore evidence about the lighting conditions—such that I am negligent in not knowing that the scene is illuminated by red lights—then my belief that the cup is red is (normatively) defeated by this unpossessed evidence.

Could defeat explain the Joyce case? Since Joyce failed to follow the rule, he may know that his experience was formed unreliably. That evidence would constitute an undercutting defeater for his belief. If Joyce's belief is defeated, then it isn't justified. Moreover, since the novice doesn't know the rule, their belief isn't defeated by this evidence. So defeat could explain why the verdict differs between the cases.

But it's unlikely that Joyce is in a position to know whether he followed the rule and thus formed his experience unreliably, since the difference between following and not following the rule is subtle. Expert perceivers do not generally attend to the precise way in which they form their experiences. Following the rule is an application of procedural knowledge, akin to following a rule of inference. When we do the latter, we do not reason *from* the rule, we simply reason *in accordance* with it. Similarly, when Joyce follows the rule, he merely perceives in accordance with it. A consequence is that Joyce can irresponsibly fail to follow the rule—just as one might irresponsibly fail to follow a rule of inference—without knowing that he failed to follow it. If Joyce doesn't know whether he failed to follow the rule, then that knowledge cannot constitute a classical defeater for his belief.

Could Joyce's unjustified belief be explained by a normative defeater? That is, *should* Joyce know that he failed to follow the rule? The question is not simply whether such a case is possible, since of course it is, but whether all cases of irresponsible experience formation entail the existence of a normative defeater. That would be true only if an obligation to follow a procedural rule entails an obligation to know whether you've succeeded in following that rule. But there is no general obligation to know what procedures you've employed in arriving at a belief or experience. In the belief case, that would entail the clearly false claim that higher-order defeaters accompany all irrational inferences. In the experience case, it would entail an obligation to attend to the application of habitual perceptual rules, even when doing so decreases performance (as it often does). That would be a surprising requirement. If there isn't that requirement, then it's possible to be obligated to follow a rule about how to form one's experience without an obligation to know whether you've succeeded. In such cases, irresponsible experience formation will not entail the existence of a normative defeater.

In summary, while it's true that evidence of Joyce's failure to follow the rule would constitute a defeater for his belief, there's no reason to think he either possesses or should possess this evidence. More generally, whenever there's an obligation to follow a rule about the formation of one's experience without an obligation to know whether one succeeds in following that rule, irresponsible experience formation, but not defeat, can explain an unjustified perceptual belief.

4.3.2 Reliability

While internalists typically appeal to defeat to explain a change in the doxastic justification of a perceptual belief, externalists can explain the change more directly by appeal to the experience's unreliability. By hypothesis, the rule that Joyce failed to follow is one that improves the reliability of his experiences. If doxastic justification is at least partially determined by the reliability of the process that produces a belief, then failing to follow the rule may result in a belief that is less justified than a belief formed in adherence to the rule. This could explain why Joyce's belief is unjustified.

The proposal faces two immediate problems. First, the fact that following the rule improves reliability does not entail that failing to follow the rule drops reliability below the threshold for *prima facie* justification. Having failed to follow the rule, Joyce uses ordinary perceptual procedures in forming his judgment about the play. Reliabilists typically assume that these procedures are reliable enough for *prima facie* justification. So it's not clear how the standard reliabilist picture can explain Joyce's unjustified belief.

But suppose it does. Then the explanation faces a second problem: the novice and Joyce employed the same belief-forming process in the same context. It follows that they were equally reliable. Either this reliability is above the threshold—and both are justified—or it's below the threshold—and neither is justified. Thus, reliability cannot explain why Joyce's belief, *but not the novice's*, is unjustified.

[Removed for review] has suggested another possible response on behalf of the reliabilist. On the 'alternative reliable process' (ARP) account of defeat, S's belief that P is defeated if S has available a more reliable process which would result in S's not believing that P (Goldman, 1979; Lyons, 2009). Intuitively, Joyce, but not the novice, has available to him the more reliable process incorporating application of the rule. And by hypothesis applying that rule would have resulted in the correct ruling that Donald was not safe. So, if we assume that experience formation is a partial constituent of the relevant process that produces a belief, it follows from ARP that Joyce's belief, but not the novice's, is defeated and thus unjustified.

This is not what Goldman or Lyons intended in extending ARP as a reliabilist theory of defeat. As I've just argued, Joyce's unjustified belief is not explainable by the presence of a defeater. And if it were, the relevant defeater would be *evidence* that Joyce formed his experience irresponsibly, not the *mere fact* that he did. Goldman (1979, 20) makes clear that he doesn't intend ARP to include application of alternative evidence-gathering processes like Joyce's attending differently to the force play:

> "[I]t seems implausible to say all 'available' processes ought to be used, at least if we include such processes as gathering new evidence. Surely a belief can sometimes be justified even if additional evidence-gathering would yield a different doxastic attitude. What I think we should have in mind here are such additional processes as calling previously acquired evidence to mind, assessing the implications of that evidence, etc."

It's interesting, then, that applying ARP to the Joyce case gets the case right. Perhaps Goldman was too quick to dismiss broader application of ARP.

However, ARP fails to capture a variant of the Joyce case. In Section 4.2, I argued that perceptual experts like Joyce should be held to higher epistemic standards. It follows that Joyce is unjustified in failing to follow the rule even if, either way, he comes to the correct conclusion that Donald was out. But since ARP is a theory of defeat, it requires that the alternative reliable process result in a different doxastic attitude. So if Joyce would form the same belief regardless of whether he follows the rule, irresponsible experience formation, but not ARP, can explain why his belief is unjustified.

What if we remove ARP's requirement that the alternative process result in a different belief? Then the view says that S's belief in P is unjustified if there is a more reliable process that S could have employed which also results in a belief whether-P. That principle would explain both the original and modified Joyce cases. It would do so by entailing that one is justified in a perceptual belief only if one employs the most reliable available process in forming that belief. This, in turn, entails that the subject should use any available procedures to form the most apt experience with respect to that belief. That is, the view entails Downgrade. It thus no longer provides a *competing* explanation of the Joyce case.

In summary, mere reliability cannot explain why Joyce's belief, but not the novice's, is unjustified. A reliabilist account of defeat—but not the version typically defended by reliabilists—can account for the cases, but not the modified cases in which following and not following the rule result in the same belief. There is a version of reliabilism that can explain all of the cases, but that version is also committed to Downgrade. I conclude that reliabilist resources do not provide a competing explanation of Joyce-like cases.

4.3.3 Inference

Susanna Siegel (2017) has recently offered a powerful defense of the view that perceptual experiences are caused by inferences whose rational character influences their epistemic status. When an experience is caused by an irrational inference, she argues, its justificatory force is downgraded. Thus, if it can be shown that Joyce's experience was formed by an irrational inference, Siegel's account can explain the verdict that Joyce's belief is unjustified without appeal to irresponsible experience formation.

It's instructive when looking for such an inference that Siegel's paradigm cases look very different from the Joyce case. Whatever the exact nature of Joyce's failure, it's related to the fact that he failed to follow a rule about how to place himself physically and mentally in order to best perceive force plays. In Siegel's cases, by contrast, the subject is apparently irrationally influenced by *some other mental state.* Consider:

> Anger: Before seeing Jack, Jill fears that Jack is angry at her. When she sees him, her fear causes her to have a visual experience in which he looks angry to her. (Siegel 2013)

Siegel claims that Jill's experience is downgraded because of an irrational influence of her fear on her experience. That influence amounts to an irrational inference from Jill's 'outlook'—the sum-total of her beliefs, desires, assumptions, etc. regarding Jack and anger—to her experience.

This explanation of what goes wrong in Anger requires that one accept four controversial claims. The first two hang together: since Siegel's explanation involves an inference from a fear to an experience with the content 'Jack is angry at me', it requires, first, that experiences have content and, second, that some of that content is rich. While Siegel has argued extensively for these claims, they remain the subject of much debate.

The third controversial claim required for Siegel's account is that perceptual contents are cognitively penetrable. On her view, experiences are downgraded when they are formed by an irrational inference from a subject's outlook—her beliefs, assumptions, moods, and so forth—to her experience. This is not an accident: Siegel's argument that experiences are formed by inferences, discussed below, is that the influence of cognitive states on experiences is epistemically on par with their influence on beliefs. Thus, the inferential character of perception is tied to its cognitive penetrability. But it's not clear whether experiences are penetrable in this way.

Finally, of course, Siegel's view requires that this cognitive penetration amounts to a rationally-assessable inference. This too is controversial. After all, paradigm rational inferences do not look like cognitive penetration. Logical reasoning, for example, is typically conscious, it involves recognizing that a conclusion is supported by some premises, and our drawing the conclusion is optional. Cognitive penetration, by contrast, is never conscious, it never involves recognizing that an experience is supported by some premises, and perceptual experiences are paradigmatically non-optional. Siegel has responded to each of these points, showing that there are inferentially formed beliefs that lack each of these features (2017, Ch. 5). And she presents a positive case that experiences are formed by inferences. She points out that cases like Anger are exactly analogous to inferentially irrational beliefs. Jill's fear influences her experiences in a way that, were the result a belief, we'd regard as inferentially irrational. Her conclusion gives 'improper weight' to her fears (2017, 5):

> Her fear makes her respond to the blank stare by experiencing Jack as angry, just as in a case of fearful thinking, it might make her respond to the blank stare by believing that Jack is angry. Here, experiences are responses to other experiences, just as conclusions are responses to inferential inputs. And like inferential responses generally, responses to experiences can be explained by the subject's rational sensitivity to the relations, as the subject sees them, between Jack's blank stare and the rest of her outlook concerning Jack and anger (2017, 118).

Since fearful seeing and fearful thinking go wrong in the same way, and since we say that fearful thinking results from an irrational inference, we should say the same of fearful seeing.

But the claim remains controversial. In particular, many have argued that experiences, unlike beliefs, cannot be epistemically 'based' on their causes and thus cannot by formed by inferences (Lyons, 2011, 2016; Ghijsen, 2015, 2016, 2018; Beck *et. al,* 2018).

It's not my goal to relitigate these disputes. My point is simply that Siegel's account of downgrade depends on each of these commitments; and that irrational inference is a competing explanation in Joyce-like cases only if these commitments turn out to be true. Irresponsible belief formation, by contrast, entails none of these commitments. It's neutral on our account of perceptual content, so long as one accepts that some perceptual experiences are more apt for a given belief than others. It's neutral about which kinds of influences—cognitive penetration, attention, and so forth—are the ones for which we're epistemically responsible, so long as we can regulate those influences with our actions. And it doesn't require any particular rational connection between an action and an experience, so long as the action meets the conditions for responsibility.

We're now in a position to see why an inferential explanation of the Joyce cases is not available. Such an explanation would require (1) that Joyce's experience be detrimentally influenced by his outlook, (2) that this influence is structurally similar to inferentially irrational belief-formation, and (3) that this influence not obtain in the novice case. None of these features are present in the Joyce cases. There's no reason to think—and anyway the case doesn't require that we think—that Joyce antecedently believed, desired, feared, or had any particular attitude toward Donald's being safe. The judgment that Joyce's belief is unjustified relies on a claim about his failing to follow the rule. The difference between the good case—in which he followed that rule—and the bad case—in which he didn't—is not a difference in Joyce's outlook regarding Donald's being safe. Second, the transition from Joyce's failing to follow the rule to his experiencing Donald as safe does not resemble an inferentially irrational belief-forming process, even if that transition is itself irrational. And last, since Joyce and the novice used the same process in forming their experiences, even if Joyce's experience were formed by an irrational inference, the novice's would be as well.

One could respond that Siegel's Anger case, an apparent instance of an irrational influence on *experience* is the wrong comparison. Siegel also argues that irrational influences on *attention* influence the epistemic standing of resulting experiences (2017, 157-160). She considers, for example, a case in which an overconfident host (Vivek) attempts to determine whether his guests are enjoying themselves but, due to his presumption that they are, fails to attend to a sulky person in the corner, waiting for a ride home (2017, 168). Vivek's outlook illicitly influences his attention such that only confirming evidence is experienced. For this reason, Siegel claims that Vivek's resulting experience is irrational. Since I've claimed that Joyce's unjustified belief is explained by his failing to follow a rule about allocating his attention, it's possible that a similar explanation is available here. That would require that Joyce's failure be explainable by an irrational inference from his outlook to the allocation of his attentional resources. But without looking at exactly how that inference should go, we can see that this explanation would face the same challenges as the one considered above. Here again, since there's no reason to assume that Joyce had any particular antecedent attitude toward Donald's being safe, the case cannot be explained by an inference from such an attitude to his attention. And even if such an inferential explanation were available, it would also be available in the novice case, since they employed the same process in forming their experiences. So this explanation cannot explain why the verdicts differ between the cases.

4.4 Implications

The argument just presented relies on experts' special obligations to enact procedures not required of novices. The corresponding difference in justification, I've argued, is explainable only by the principle that beliefs based on irresponsibly formed experiences are unjustified. Thus, irresponsible experience formation has unique explanatory advantages. This argument might give the impression that the epistemic impact of irresponsible experience formation is minimal. But once the epistemic relevance of irresponsible experience formation is established, its impact redounds throughout our epistemic lives. In this final section, I'll comment briefly on this impact.

While my emphasis here has been on special epistemic obligations, it's likely that there are both general and special obligations regarding the formation of experiences. It's plausible, for example, that we have a general epistemic obligation to attend to immediately available evidence concerning occurrent beliefs. My belief that it's acceptable to feed the zoo animals, for example, is undermined by the presence, even if I don't notice it, of a sign that says otherwise (Gibbons 2006). In this way, our general epistemic obligations concerning experience formation can explain certain kinds of normative defeaters. I suggested in Section 4.2 that there are at least two *kinds* of special epistemic obligations: those grounded in what we know and those grounded in our social roles. Joyce's special obligation to follow the rule holds both in virtue of what he knows and his social role. Similar cases abound: scientists have special obligations to attend to certain kinds of evidence; police officers have special obligations to recognize when an object is (or isn't) a weapon; psychologists have special obligations to pick up on signs of trauma. Other special obligations may be contingently universal, such as an obligation, given widespread evidence, to regulate the influence of implicit biases on our experiences (Washington and Kelly 2016). In short, concerns about responsible experience formation apply throughout our epistemic lives.

Susanna Siegel's recent work has highlighted the wide variety of ways in which the etiology of experience may redound on the epistemic standing of perceptual beliefs. The responsibility view can incorporate Siegel's inferentialist framework in two ways. First, much of the literature on responsibilist accounts of doxastic justification highlight good inference as a paradigm example of responsible belief formation. Thus, if experiences are formed by inferences whose rational standing redounds on their epistemic standing, then this can be seen as a special case of the responsibility view. But I noted in the last section that Siegel's framework makes several controversial commitments not shared by the responsibility view. Perhaps most notably, the responsibility view can explain cases of epistemic downgrade without requiring that experiences be formed by inference. Thus, the responsibility view may explain Siegel's cases even if experiences are not inferentially formed. We may, for example, have general or special obligations to prevent influence of our fears on our experiences.

As discussed in Section 4.2, a direct argument for Responsibility and Downgrade, like the one developed here, establishes irresponsible experience formation as an independent explanandum for general theories of doxastic justification. If the above arguments are sound, then irresponsible experience formation is a *sui generis* influence on doxastic justification—akin to defeat—whose epistemic import must be captured by broader theories. Like Siegel's account, Downgrade challenges hard internalist positions committed to phenomenal conservatism, since it denies that, in the absence of defeaters, having an experience that p suffices for immediate justification of a belief that p. Unlike Siegel's account, however, Downgrade also challenges external accounts, since, as established in Section 4.3.2, no existing reliabilist account can capture the epistemic import of irresponsible experience formation.

Finally, the responsibility view broadens the notion of epistemic downgrade itself. The literature on epistemic downgrade has its roots in debates over the epistemic implications of cognitive penetration. If our experiences are influenced by our beliefs, desires, emotions and so forth, then it seems this should have an impact on their epistemic standing. Inferentialism is designed to vindicate this intuition. As a result, Siegel's inferentialist conception of downgrade assumes that experiences are cognitively penetrated. While the responsibility view can account for the epistemic relevance of cognitive penetration, it also explains cases in which experiences are downgraded by other causes. Joyce's failure to follow the rule, for example, is not an instance of cognitive penetration. The responsibility view thus reorients the debate over epistemic downgrade by treating cognitive penetration as a mere special case; and it extends the range of possible epistemic influence on perception to attentional effects, perceptual learning, priming, and so forth. Future investigation into downgrade should explore the epistemic impact of a broad range of possible influences on perception. More generally, the responsibility view makes room for new discussion of the relationship between our understanding of our own perceptual processing and its influences and the epistemic standing of our perceptual beliefs.

References

Adler, J. (2004). Reconciling open-mindedness and belief. *Theory and Research in Education*, *2*(2), 127-142.

Audi, R. (1986). Belief, Reason, and Inference. *Philosophical Topics* 14 (1): 27–65.

——. (2001). Doxastic Involuntarism and the Ethics of Belief. *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue,* 93-111. Oxford University Press.

Baehr, J. S. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford University Press.

Beck, O., Chirimuuta, M., Rosenhagen, R., Siegel, S., Smithies, D., & Springle, A. (2018). Discussion of Susanna Siegel's 'Can perceptual experiences be rational?' *Analytic Philosophy*, *59*(1), 175–190.

Boghossian, P. (2014). What Is Inference? *Philosophical Studies* 169 (1): 1–18.

Broome, J. (2014). Comments on Boghossian. *Philosophical Studies* 169 (1): 19–25.

Carroll, L. (1895). What the tortoise said to Achilles. *Mind*, *4*(14), 278-280.

Chudnoff, E. (2014). The Rational Roles of Intuition. *Intuitions*, 9–35. Oxford University Press.

Churchland, P. M. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, *55*(2), 167–187.

Code, L. (1987). *Epistemic Responsibility*. Brown University Press.

Cohen, J. (1977). *The Probable and the Provable.* Oxford University Press.

Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3), 667–695.

Firestone, C., & Scholl, B. (2015). Cognition Does Not Affect Perception: Evaluating the Evidence for 'Top-Down' Effects. *Behavioral and Brain Sciences*.

Fodor, J. (1988). A Reply to Churchland's" Perceptual Plasticity and Theoretical Neutrality". *Philosophy of Science*, *55*(2), 188–198.

Ghijsen, H. (2015). Grounding Perceptual Dogmatism: What are Perceptual Seemings? *The Southern Journal of Philosophy*, *53*(2), 196–215.

——. (2016). The real epistemic problem of cognitive penetration. *Philosophical Studies*, *173*(6), 1457–1475.

——. (2018). How to Explain the Rationality of Perception. *Analysis*, *78*(3), 500–512.

Gibbons, J. (2006). Access externalism. *Mind*, *115*(457), 19–39.

Goldman, A. (1979). What is justified belief?. *Justification and knowledge*, 1-23. Springer.

Helton, G. (2016). Recent issues in high-level perception: High-Level Perception. *Philosophy Compass*, *11*(12), 851–862.

Kornblith, H. (1983). Justified Belief and Epistemically Responsible Action. *The Philosophical Review*, *92*(1), 33–48.

Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.

——. (2011). Assertion and Isolated Second-Hand Knowledge. *Assertion: New Philosophical Essays*. Oxford University Press.

Littlejohn, C. (2017). Truth, knowledge, and the standard of proof in criminal law. *Synthese*, 1-34.

Lyons, J. (2009). *Perception and basic beliefs: Zombies, modules, and the problem of the external world.* Oxford University Press.

——. (2011). Circularity, reliability, and the cognitive penetrability of perception. *Philosophical Issues*, *21*(1), 289–311.

——. (2016). Inferentialism and cognitive penetration of perception. *Episteme*, *13*(1), 1–28.

McGrath, M. (2013). Phenomenal conservatism and cognitive penetration: The 'bad basis' counterexamples. *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, 225–247. Oxford University Press.

Montmarquet, J. (1993). *Epistemic virtue and doxastic responsibility*. Lanham, MD: Rowman & Littlefield.

Peels, R. (2017). *Responsible belief: a theory in ethics and epistemology*. Oxford University Press.

Pollock, J., & Cruz, J. (1999). *Contemporary Theories of Knowledge*. Rowman & Littlefield Publishers.

Siegel, S. (2011). *The contents of visual experience*. Oxford University Press.

——. (2013). The epistemic impact of the etiology of experience. *Philosophical Studies*, *162*(3), 697-722.

——. (2017). *The rationality of perception*. Oxford University Press.

Sosa, E. (2007). *A virtue epistemology*. Clarendon Press ; Oxford University Press.

Staffel, J. (2016). Beliefs, buses and lotteries: Why rational belief can't be stably high credence. *Philosophical Studies*, *173*(7), 1721-1734.

Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, 49(3): 199–219.

Washington, N., & Kelly, D. (2016). Who's Responsible for This?. *Moral Responsibility, Externalism, and Knowledge about Implicit Bias, Implicit Bias and Philosophy*, 2, 12-36.

Ye, R. (2018). Higher-order defeat and intellectual responsibility. *Synthese*, 1-21.

Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.

Chapter 5: Degree Encroachment

This chapter articulates a new theory of pragmatic or moral encroachment according to which traditionally non-epistemic factors influence whether a subject has knowledge by helping to determine which epistemic reasons and procedures they employ in forming their belief. Since doxastic justification is a function of reasons and procedures, this view entails that non-epistemic factors encroach on the epistemic by influencing the *degree* of graded justification the subject has for her belief.

Section 5.1 discusses the extant literature on encroachment and distinguishes between 'degree encroachment' and 'threshold encroachment'. Section 5.2 argues for degree encroachment on reasons. Section 5.3 argues for degree encroachment on procedures. Section 5.4 articulates two advantages of degree encroachment over its extant competitors.

5.1 Degree Encroachment and Threshold Encroachment

An epistemic theory posits encroachment only if, holding fixed traditionally epistemic factors, variations in traditionally non-epistemic factors can influence an epistemic status (such as whether a subject knows, is justified, or has a rational credence). To see this, consider the following pairs of cases:

Bank Cases: *Keith is deciding whether to deposit a paycheck on Friday or Saturday. He remembers that it was open on Saturday two weeks ago, but it's possible that there's an exception this Saturday. In Low-Stakes, it makes no difference whether Keith deposits the check this weekend. In High-Stakes, a very important check will bounce if the paycheck is not deposited this weekend. In Low-Stakes, but not in*

High-Stakes, Keith knows the bank is open on Saturday. (simplified from DeRose, 2005)

Street Cases: *Sarah is deciding whether to cross the street in response to a perceived threat. In Low-Stakes, she judges that a pit-bull on her side of the street is more likely to bite her than any dog across the street. In High-Stakes, she judges that a black person on her side of the street is more likely to steal her purse than anyone across the street. Her evidence for these claims is equally strong. Nonetheless, in Low-Stakes, but not in High-Stakes, Sarah has knowledge.* (simplified from Moss, 2018)

Both pairs of cases conclude with a commonly-held judgment about the differences between the cases. The claim is that the difference in stakes results in a difference in whether the subject knows, even though traditional epistemic factors are the same. The first case is an instance of pragmatic encroachment because the subject's practical stakes intuitively influence whether the subject knows. The second case is an instance of moral encroachment because the subject's moral stakes intuitively influence whether the subject knows. The general phenomenon of traditionally non-epistemic factors influencing an epistemic status is sometimes called 'pragmatic encroachment,' with moral encroachment as a special case. I find this language misleading, since on most meta-ethical theories moral stakes do not necessarily influence practical stakes. I'll refer to the general phenomenon simply as 'encroachment' with pragmatic and moral encroachment as special cases. Nothing will turn on this choice of language. While I will not defend either pragmatic or moral encroachment in particular, my theory can explain instances of both.

Brian Weatherson (2005) has made an important clarification to our concept of encroachment. A non-epistemic influence on an epistemic status counts as encroachment only if it does so by influencing *a normative component* of that status. A belief, by contrast, is a non-normative component of knowledge. His 2005 article defends (but he now rejects; see his 2012) the view that purported cases of encroachment on knowledge are merely influences of non-epistemic factors on belief.

Most theories of encroachment (including my own) aim to explain influences of non-epistemic factors on knowledge by positing influence of those factors on knowledge-level justification. I'll assume, with the literature, that justification above some threshold is a necessary condition on knowledge. Thus, non-epistemic factors can influence whether a subject has knowledge-level justification either by influencing the location of that threshold—what I'll call 'threshold encroachment'—or by influencing the degree to which the subject's belief is justified—what I'll call 'degree encroachment'. Most defenders of encroachment defend threshold encroachment. My view is a kind of degree encroachment.

5.2 Encroachment on Reasons

Alex Worsnip (2021) suggests we think about justification in terms of a 'beaker of reasons' (see Fig. 5.1). The more reasons one has for one's belief, the fuller the beaker. If one has enough reasons, they pass a threshold and the belief is knowledge-level justified. Non-epistemic factors can influence

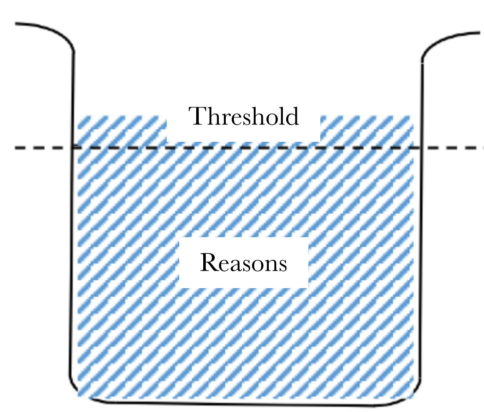

Figure 5.1: The 'Beaker of Reasons'. Based on a figure from Worsnip (2021)

the beaker by changing either which reasons are in the beaker (a kind of degree encroachment) or the location of the threshold (a kind of threshold encroachment).

Degree encroachment has not been widely discussed in the literature. The only version discussed, to my knowledge, is the view that non-epistemic factors influence the contents of the beaker by *constituting* reasons for and against belief. I call this the 'reasons-constituting' view. A reasons-constituting explanation of the bank cases, for example, would say that, in High-Stakes, the fact that it matters to Keith whether the bank is open is a reason against believing that the bank is open on Saturday, lowering the level in the beaker below the threshold. This view is widely rejected in the literature on the grounds that it faces a devastating 'wrong kind of reasons' objection: even if we think that stakes influence justification, stakes are not evidentially related to whether the bank is open on Saturday. It's therefore hard to see how they could constitute *epistemic* reasons for or against believing the bank is open on Saturday (Worsnip, 2021; Fritz, 2020; Bolinger, 2020).

There's another version of degree encroachment that hasn't been discussed in the literature, however. That's the view that non-epistemic factors help determine which epistemic reasons one possesses without themselves constituting reasons for or against belief. I call this the 'reasons-determining' view. Like the reasons-constituting view, the reasons-determining view says that non-epistemic factors influence the contents of the beaker. Thus, the reasons-determining view is also a kind of degree encroachment. But unlike the reasons-constituting view, the reasons-determining view doesn't claim that non-epistemic factors are themselves contents of the beaker of reasons. Instead, non-epistemic factors help determine which classically epistemic reasons enter the beaker.

How could non-epistemic factors help determine which epistemic reasons enter the beaker? Suppose that I believe that my flight will leave on time, but, unbeknownst to me, the flight board (which is right in front of me) says that my flight is delayed. The traditional view is that a piece of evidence is relevant to justification only if it is accessible upon reflection (Pappas, 2017). The contents of the board are not accessible to me upon reflection, so this view entails that it cannot be a reason for me to abandon my belief. Since it's not a reason for me to abandon my belief, it cannot undermine my justification.

But this traditional view is challenged by cases in which I clearly *should* have access to a piece of evidence that undermines my belief. Suppose, for example, that I'm supposed to pick up my child from daycare on the other end of this flight. I'm obligated to make arrangements if I won't be making the pickup on time. Then, plausibly, I *should* check the board (especially since it's right in front of me) and so I *should* have access to the fact that my flight is delayed. I cannot justifiably believe that P if I should have conclusive evidence that not-P. So I need not have reflective access to evidence for it to count as a reason against my belief (Gibbons, 2006; Lackey, 2008).

One way, then, that non-epistemic factors can determine which epistemic reasons enter the beaker is if non-epistemic factors can make it the case that I should have certain evidence. To see this, contrast the case just mentioned with an exactly analogous case in which it doesn't particularly matter (morally or practically) whether my flight lands on time. In this case, I plausibly don't have an obligation to check the board and so can justifiably believe that my flight will land on time. The difference between the cases is the stakes; and the result of this difference is a difference in which reasons enter the beaker.

Notice that the reason that I possess in the high stakes case, but not in the low stakes case—the evidence from the flight board—is a straightforwardly epistemic reason. It's truth or falsity is probabilistically relevant to the truth or falsity of my belief, for example. Once we stipulate that this evidence is in my beaker of reasons, the conclusion that my belief is unjustified follows from ordinary epistemic principles. The stakes, then, are not themselves reasons for or against belief, but they do help determine which epistemic reasons the subject has. Also notice that the argument that I'm unjustified in high stakes, but not in low stakes, does not depend on a prior judgment about whether it's appropriate to attribute knowledge to me. Instead, the argument turns on independently-motivated arguments about the conditions for reason possession. These two features—allowing for a traditional function from epistemic factors to justification, and admitting of independent motivation from outside of the encroachment literature—are core strengths of my view.

The reasons-determining view is similar to a 'relevant alternatives' contextualism according to which non-epistemic factors help determine which alternatives a subject must rule out to count as knowing (e.g. Moss, 2018). They agree, for example, that my knowledge is undermined in the high-stakes flight case because I failed to check the board. And they give similar explanations of why: high-stakes make it the case that I should check whether my flight is delayed. But the views differ in two ways. First, the relevant alternatives view is, as it's typically formulated, a version of threshold encroachment, while the reasons-determining view is a version of degree encroachment. The reasons-determining view says that higher stakes require that I respond to a different set of epistemic reasons than lower stakes. Thus, my belief is less justified in higher-stakes situations. The relevant alternatives view, by contrast, says that higher stakes require that I respond to the same reasons in a different way to count as justified. Higher stakes raise the threshold for knowledge-level justification. As a result, the relevant alternatives view—and all threshold accounts—have the counterintuitive result that responding to the same set of epistemic reasons in the same way does not determine whether a belief is justified across all cases. My view, by contrast, does not have this entailment.

A second difference between the views is their scope. The relevant alternatives view says that a belief that P doesn't count as knowledge unless the subject rules out all the relevant alternatives to P. That is, believing that P is justified only if it's epistemically impossible (relative to the relevant set of possibilities) that not-P. The corresponding reasons-determining view would be that believing that P isn't justified if one possesses (or should possess) reasons to believe that not-P. A reason to believe that not-P is a rebutting defeater for a belief that P. But not all defeaters are rebutting defeaters. And not all reasons are defeaters. So, the reasons-determining view allows for more ways in which non-epistemic factors can influence the justification status of a belief. For example, pragmatic factors may make it the case that one has a (normative) undermining defeater for one's belief. Perhaps in high-stakes scenarios, for example, I should rule out the possibility that the lighting conditions are influencing my visual judgments. This consideration is independent of whether my visual judgments are incorrect. Or suppose that, in a high-stakes scenario, the flight board confirms my belief that I'll make my flight, but I nonetheless fail to check it. The evidence I *should* respond to (including the content of the board) might propositionally justify my belief, but I might fail to appropriately base my belief on this evidence. In this case, raised stakes result in my belief lacking justification, not because of the presence of a defeater, but because my belief fails to be based on the appropriate reasons.

5.3 Encroachment on Procedures

The last section employed Worsnip's 'beaker of reasons' framework and argued that there is a sensible way to defend degree encroachment on reasons: rather than non-epistemic factors *constituting* reasons for or against beliefs (and thus entering the beaker of reasons), they instead help *determine* which epistemic reasons one has. If so, then they can explain why one's reasons fall below the threshold in high-stakes scenarios without appealing to a change in that threshold.

I must now make an important clarification. Propositional justification concerns whether a subject has sufficient reason to believe a given proposition. Doxastic justification concerns whether a belief is appropriately held. It's plausible that one is knowledge-level *propositionally* justified whenever one's reasons for a target proposition rise above a threshold. The beaker of reasons is thus a useful framework for thinking about propositional justification. But it is not plausible that one is knowledge-level *doxastically* justified whenever one's reasons rise above the threshold. Doxastic justification also requires that one execute the right procedures, that is, that one respond to one's reasons *in the right way* (Swain, 1979; Pollock & Cruz, 1999; Korcz, 2000; Feldman, 2002; Kvanvig, 2003). So the beaker of reasons is not a good framework for thinking about doxastic justification.

The distinction between propositional and doxastic justification is not widely discussed in the encroachment literature. If you accept the threshold view, it doesn't particularly matter what the determinants of justification are, so long as one accepts that raised stakes raise the threshold. In that context, assuming that justification is only determined by one's reasons is an innocent simplifying assumption. By contrast, if you accept the degree view, as I do, it's important that you clarify exactly which determinants of justification are influenced by non-epistemic factors, and how. I've just explored how non-epistemic factors might influence which reasons we have (and thus propositional justification). It's therefore worth exploring whether non-epistemic factors independently influence doxastic justification, that is, whether non-epistemic factors influence what procedures we employ in arriving at our beliefs.

Let's call the view that pragmatic factors help determine which procedures one employs in arriving at a belief the 'procedures-determining' view. You might find the procedures-determining view implausible if you think that the procedures which determine doxastic justification are constituted by purely descriptive facts. Surely my practical or moral situation cannot change whether I, for example, am guilty of employing a logical fallacy. But epistemically relevant features of our procedures—and in particular *negative* features of our procedures—are influenced by normative facts. Consider, for example, whether someone has ignored relevant evidence. That factor—which is clearly an epistemically relevant feature of the procedure by which they formed their belief—is at least partially normative in nature. One counts as ignoring evidence only if there is evidence one *should have* attended to. And, if the reasons-determining view is right, whether one should have attended to certain evidence (i.e., whether it counts as a reason one should respond to) is partially determined by one's practical or moral situation. Thus, the procedures-determining view follows from the reasons-determining view.

But could there be influences on our procedures which are not also influences on our epistemic reasons? Plausibly. Consider two situations in which you perform a simple arithmetical calculation without double checking to ensure you've done it correctly. In a Low-Stakes scenario, a minor error in the calculation would have little impact on anyone. In a High-Stakes scenario, an error could have devastating effects. In both cases, you have the same set of epistemic reasons (the numbers are the same, knowledge of procedures of arithmetic are the same, etc.). But it's plausible that you should have double checked in the high-stakes scenario. If so, then your failure to do so may impact whether you are doxastically justified. This, then, would be an influence of a non-epistemic factor on one's procedures independently of any influence on one's reasons.

Notice that the procedures-determining view has similar features as the reasons-determining one. They collectively allow for a function from the classical determinants of justification to knowledge-level justification. And they are both motivated without reference to intuitions about whether it is appropriate to attribute knowledge in particular cases. I'll now argue that each is a crucial advantage of my view over the extant threshold accounts of encroachment.

5.4. Two Advantages of Degree Encroachment

In the preceding sections, I've aimed to show that degree encroachment is a viable account of how non-epistemic factors might influence whether a subject is knowledge-level justified. The possibility of encroachment on graded justification has been overlooked, I think, because it has been falsely assumed that non-epistemic factors can influence graded justification only by *constituting* reasons for or against a belief. Once the distinction between a factor constituting a reason and merely helping to determine it is acknowledged, degree encroachment becomes a more plausible account. In what remains of this chapter, I aim to show that, in addition to being viable, degree encroachment has two crucial advantages over the threshold mechanism: it admits of independent motivation beyond the encroachment literature, and it allows for a more traditional relationship between epistemic factors and justification.

5.4.1 Independent Motivation for Degree Encroachment

Most accounts of encroachment are motivated by their ability to explain cases in which knowledge apparently depends on a non-epistemic factor. DeRose's original account (2005), for example, relied on our accepting the judgments given in the bank cases. Similarly, recent accounts of moral encroachment (e.g., Moss, 2018) rely on judgments about whether it is appropriate to attribute knowledge in cases of high moral significance (such as the street cases). Much of the literature, therefore, has revolved around whether the verdicts in these cases are correct, and whether they might be explained without appeal to encroachment (e.g., Weatherson, 2005).

Like its opponents, my view can explain apparent cases of pragmatic or moral encroachment on knowledge. In the bank cases, for example, Keith's high stakes apparently require him to enact additional procedures (e.g., calling up the bank) before forming a belief about the bank's hours. Since he fails to enact this procedure, he is less justified than if he did. Moreover, he plausibly lacks this obligation in the low-stakes case. Thus, degree encroachment can explain our different verdicts across the cases.

But since most accounts of encroachment can offer *some* explanation of the plausible cases of encroachment on knowledge, independent grounds are required for adjudicating between them. Much of the literature, therefore, concerns theoretical reasons that one account might be better than the other. Much of this debate does not leave the encroachment literature. My view, by contrast, is independently motivated by arguments for normative conditions for reasons and procedures. The arguments given above for the reasons-determining view, for example, are generalizations from arguments for so-called ‘normative defeaters’ (Lackey, 2008). More generally, debates about the conditions under which we possess a reason or enact a particular procedure largely take place outside of the encroachment literature. A strength of my view, then, is that it can explain intuitions both within and outside of the encroachment literature.

4.2. Epistemic Reasons and Justification

The second advantage of my view concerns the relationship between epistemic reasons and knowledge. I mentioned above that the reasons-constituting view is widely regarded as facing a devastating ‘wrong kinds of reasons’ objection. Pragmatic reasons, for example, simply do not seem to be the right kinds of reasons to influence justification. Threshold accounts face a similar objection. Michael Blome-Tillmann (2009) has argued, for example, that threshold accounts unacceptably entail the truth of claims like, “S knows that p, but if it were more important, she wouldn’t know”.

These considerations can be generalized into the following general objection to accounts of encroachment. Since knowledge-level justification is an *epistemic* status, whether a subject has it should be a function of intuitively epistemic factors such as evidence, procedures, and so forth. We should be able to stipulate all of the epistemic factors and know whether a subject is justified in believing as they do. But if any of the extant accounts of encroachment are right, then there can be no such function. Knowing the epistemic factors *never* determines whether the subject knows, since even in the limit case in which their evidence deductively entails their conclusion, there may or may not be an additional procedure (such as double-checking their reasoning) that is required for knowledge. This is a powerful consideration against extant accounts of encroachment.

But notice that since, on my view, non-epistemic factors help to determine which reasons and procedures a subject employs in arriving at her belief without influencing how these factors relate to knowledge-level justification, my view is consistent with the existence of a function from non-epistemic factors to an epistemic status like knowledge-level justification. My view is thus unique in positing encroachment while allowing for a traditional relationship between epistemic factors and statuses.

References

Blome-Tillmann, M. 2009. Contextualism, Subject-Sensitive Invariantism, and the Interaction of 'Knowledge'-Ascriptions with Modal and Temporal Operators. *Philosophy and Phenomenological Research*, 79(2): 315–331. doi:10.1111/j.1933-1592.2009.00280.x

Bolinger, R. J. (2020). Varieties of moral encroachment. *Philosophical Perspectives*, *34*(1), 5-26.

DeRose, K. 2005. The ordinary language basis for contextualism, and the new invariantism. *The Philosophical Quarterly*, *55*(219), 172-198.

Feldman, R. (2002). *Epistemology*. Upper Saddle River, NJ: Prentice Hall

Fritz, J. (2020). Moral encroachment and reasons of the wrong kind. *Philosophical Studies*, *177*(10), 3051-3070.

Gibbons, J. 2006. Access externalism. *Mind*, *115*(457), 19–39.

Korcz, K. A. (2000). The causal-doxastic theory of the basing relation. *Canadian Journal of Philosophy*, *30*(4), 525-550.

Kvanvig, J. L. (2003). Propositionalism and the perspectival character of justification. *American Philosophical Quarterly*, *40*(1), 3-17.

Lackey, J. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.

Moss, S. 2018. IX—Moral Encroachment. *Proceedings of the Aristotelian Society*, *118*(2), 177–205. https://doi.org/10.1093/arisoc/aoy007

Pappas, George. 2017."Internalist vs. Externalist Conceptions of Epistemic Justification", The Stanford Encyclopedia of Philosophy (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2017/entries/justep-intext/>.

Pollock, J. L., & Cruz, J. (1999). *Contemporary theories of knowledge* (Vol. 35). Rowman & Littlefield.

Swain, M. (1979). Justification and the Basis of Belief. In *Justification and knowledge* (pp. 25-49). Springer, Dordrecht.

Turri, J. 2010. On the Relationship between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research*, *80*(2), 312–326. https://doi.org/10.1111/j.1933-1592.2010.00331.x

Weatherson, B. 2005. Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*, *19*(1), 417–443. https://doi.org/10.1111/j.1520-8583.2005.00068.x

Weatherson, B. 2012. Knowledge, Bets, and Interests. In J. Brown & M. Gerken (Eds.), *Knowledge Ascriptions* (pp. 75–103). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199693702.003.0004

Worsnip, A. (2021). Can pragmatists be moderate?. *Philosophy and Phenomenological Research*, *102*(3), 531-558.