

5-2020

Perceptual Characterization: On Perceptual Learning and Perspectival Sedimentation

Anthony Holdier
University of Arkansas, Fayetteville

Follow this and additional works at: <https://scholarworks.uark.edu/etd>



Part of the [Cognitive Psychology Commons](#), and the [Epistemology Commons](#)

Citation

Holdier, A. (2020). Perceptual Characterization: On Perceptual Learning and Perspectival Sedimentation. *Graduate Theses and Dissertations* Retrieved from <https://scholarworks.uark.edu/etd/3656>

This Thesis is brought to you for free and open access by ScholarWorks@UARK. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of ScholarWorks@UARK. For more information, please contact scholar@uark.edu.

Perceptual Characterization: On Perceptual Learning and Perspectival Sedimentation

A thesis submitted in partial fulfillment
of the requirements for the degree of
Master of Arts in Philosophy

by

Anthony Holdier
Colorado State University
Bachelor of Arts in Anthropology, 2009
Denver Seminary
Master of Arts in Philosophy of Religion, 2012

May 2020
University of Arkansas

This thesis is approved for recommendation to the Graduate Council.

Jack Lyons, Ph.D.
Thesis Director

Amanda McMullen, Ph.D.
Committee Member

Eric Funkhouser, Ph.D.
Committee Member

Abstract

In her analysis of perspectival effects on perception, Susanna Siegel has argued that perceptual experience is directly rationally assessable and can thereby justify perceptual beliefs, save for in cases of epistemic downgrade or perceptual hijacking; I contend that the recalcitrance of known illusions poses an insurmountable problem for Siegel's thesis. In its place, I argue that a model of perceptual learning informed by the dual-aspect framework of base-level cognitive architecture proposed by Elisabeth Camp successfully answers the questions motivating Siegel's project in a manner that avoids such issues.

Dedication

For Jenni, Ellie, Audrey, and Tessa – you help me see what truly matters.

Table of Contents

I	Introduction	1
II	On Perceptual Epistemology	5
	1. Perceptual Systems Theory and MUG	6
	2. Perceptual Conservatism and MUG	8
	3. Perceptual Systems Theory, Perceptual Conservatism, PINE, and SCOPE	15
III	Why Perception is Nonrational	18
	1. Perceptual Systems as Modular Input Systems	21
	2. Perceptual Conservatism and Recalcitrant Illusions	27
	3. On Perceptual Hijacking	35
	4. Why Perception is Nonrational	38
IV	On Perspectival Epistemology	41
	1. Characterizations and Concepts	44
	a. What Characterizations Are	44
	b. What Characterizations Are Not	49
	c. The Camp Doctrine	52
	2. Perceptual Characterization and the Contents of Perception	54
V	Perspectival Sedimentation	63
	References	66

List of Figures

1	The Müller-Lyer Illusion	18
2	The Necker Cube	23
3	The Duck-Rabbit	23
4	“My Wife and My Mother-in-Law”	48

Epigraph

All knowledge takes its place within the horizons opened up by perception.

- Maurice Merleau-Ponty ([1945] 1962: 207)

Thoughts without concepts are empty, intuitions without concepts are blind.

- Immanuel Kant ([1781] 1998: 193-194)

I Introduction

Consider the following case:

PINE: Susan is hired to cut down all of the pine trees in a small forest containing a mixture of tree kinds. She brings Calvin to assist her, despite the fact that he has never before seen a pine tree. As they work, Susan gradually teaches Calvin how to recognize pine trees by simply pointing them out as they go. Though it takes time, pine trees eventually become visually salient to Calvin such that he can immediately recognize them: he has gained a recognitional disposition to distinguish pine trees from non-pine trees.¹

A case like this is fruitful for philosophers. Susanna Siegel (2011) initially used it to illustrate the Phenomenal Contrast argument in defense of her Rich Content View whereby high-level concepts are represented in perception. Calvin's disposition acquisition is adapted by Kevin Connolly (2019) to illustrate his Offloading model of perceptual learning; Stokes (2018) uses it to discuss the perception of aesthetic properties; Macpherson (2012) advances a similar argument regarding the cognitive penetration of perception.

Consider now a second case:

SCOPE: Thomas is a well-informed biologist in the early seventeenth century who believes in preformationism (the theory that organisms develop from smaller versions of themselves). When Thomas inspects a sample with his microscope, he has a perceptual experience of a tiny embryo inside of a sperm cell.²

Because perception (as opposed to hallucination or illusion) is factive and preformationism is false, it is impossible for Thomas to perceive an embryo inside of a sperm cell. Nevertheless, his perceptual experience might be such that, in virtue of having the experience, Thomas receives confirmation of his (incorrect) pre-theoretical position — in Siegel's terms, his "prior outlook" — via perception.

¹ This case is heavily based on the precursor to the Phenomenal Contrast Argument found in Siegel (2011: 100).

² This case is also based on an example from Siegel (2017: 7).

The cognitive penetration of a system is “when a system takes beliefs, desires, goals, or other similar person-level states as input, where ‘input’ can be direct or indirect, but cannot be behaviorally mediated” (Lyons 2015a: 110). Notably, this definition delimits cognitive penetration from any kind of top-down influence whatsoever: as Brogaard and Chomanski (2015) point out, higher-level brain states that control a subject’s attention or intention can indeed *effect* an agent’s perceptual experience (by, say, directing how a person chooses to foveate), but they do not *penetrate* perception (471): purported cases of penetrated experience have identical behavioral etiologies with non-penetrated experiences. In Siegel’s treatment of a case like PINE, she argues that as Calvin forms beliefs about pine trees, those beliefs penetrate his perceptual experience, thereby altering Calvin’s visual phenomenology and what Calvin is *prima facie* justified to believe (2012: 204–205). To Siegel, the synchronic effects of Calvin’s doxastic states are sufficient to provoke a contrast in the content of his visual phenomenal states (or how the forest looks to him): what Calvin believes about pine trees affects what he sees.³

However, the slow, laborious process required to manifest the change in Calvin’s visual experience suggests that PINE should instead be treated *not* as a case of cognitive penetration, but as a case of perceptual learning, whereby Calvin trains his perceptual systems to render different outputs such that he gradually associates <pine tree> with the (unchanging) contents of his visual experience and thereby gradually gains the disposition to see pine trees *as pine trees*: the property of “pine tree” is eventually encoded directly by his perceptual systems (rather than secondarily as a matter of cognition).⁴ Just as cognitive penetration must be disambiguated from any simple top-down synchronic influences on experience (such as attention), it must also be

³ Although the thesis that perception has contents is itself controversial, the focus of this work is to consider the ramifications of a particular version of a content thesis; consequently, anti-representational or other models of perceptions that deny perceptual contents will be functionally disregarded. For a contrary view, see Johnston (2014, 2006).

⁴ Throughout this work, I shall use the typography “<...x...>” to denote “the concept of x.”

distinguished from the diachronic training of a subject's perceptual systems (Connolly 2014: 1408); whereas cognitive penetration threatens to undermine the epistemological foundations of perceptual processes in a host of ways,⁵ the products of perceptual learning are less transient and, thereby, more reliable (Gatzia and Brogaard 2017: 2).⁶ Given that Calvin's ability to see pine trees as pine trees is provoked by his experience of actual pine trees, this process of perceptual learning does not negatively affect Calvin's overall rational standing — indeed, it enables him to reliably form a large number of justified beliefs regarding pine trees in his vicinity.

On one hand, the SCOPE case is simple to interpret because it is similar to PINE: given that Thomas has (consciously or otherwise) trained himself to view the world from a preformationist perspective, he has a disposition to see things “preformationist-ly” and this includes things like embryos in sperm cells. Consequently, Thomas' perceptual experience has been affected by his perspective such that he believes himself to have observed an embryo on his sample slide (much like how Calvin's perceptual experience has been affected by his newfound perspective as a tree-perceiver such that, upon seeing a pine tree, he believes himself to have observed a pine tree). On the other hand, there is no embryo for Thomas to see because sperm cells do not actually contain such things; the influence of his perspective on his perceptual experience has negatively impacted Thomas' rational standing.

Thomas' situation is roughly what Siegel calls the Problem of Hijacked Experience: when a subject's prior outlook overrides her perceptual inputs in a manner that influences her perceptual judgments, Siegel says that the outlook *hijacks* the subject's perceptual state (2017: 5). For many, *if* such hijacking is possible, it would simply be a case of hallucination or illusion

⁵ Fodor once remarked that the threat posed by cognitive penetration for perception is precisely that “Background beliefs, and the expectations that they engender, from time to time prove *not to be true* (1983: 38).

⁶ Siegel (2017: xix) also briefly mentions perceptual learning as “processes internal to perceptual systems,” but goes on to treat this as simply one species of the broader genus of cognitive penetration, a mistaken heuristic at least somewhat established in the literature, as will be discussed in Section III.

insofar as Thomas is forming unjustified perceptual beliefs about what he incorrectly thinks is on his microscope slide — the fact that his perceptual states are experientially faulty because of his doxastic states is ultimately not relevant for making epistemological judgments about the *phenomenology* of his visual perceptual experience (merely its epistemological explicandum); Siegel, by contrast, places a significant weight on Thomas' visual phenomenology, arguing that — although Thomas is indeed suffering from a hallucination or illusion — the phenomenology itself possesses assessable epistemic power. In order to analyze the manner and degree of experiential hijacking, Siegel has created an elaborate conceptual framework, including the *sui generis* epistemological property of epistemic charge — a valenced measurement of how one's outlook has affected one's experience.

Unfortunately, the degree to which Siegel's model centralizes the role of phenomenological experience makes her conclusions difficult to countenance in light of contemporary psychology and vision science. Her Phenomenal Contrast argument is based on an important, undefended assumption about our cognitive architecture and her Rich Content View remains challenged in the contemporary literature. Her attempts to characterize and solve the Problem of Hijacked Experience represent her engagement with such criticism, but it fails to account for a crucial counterexample.

Instead, in what follows, I will argue for two claims:

- 1) Siegel's Rationality of Perception thesis fails to explain the recalcitrance of perceptual illusions: it is unsound.
- 2) An expanded model of human cognitive architecture — namely, that defended by Elisabeth Camp — provides a more powerful explanation for perspectival influences on perception.

In so doing, I aim to not only offer a corrective to an important theory in the contemporary philosophy of perception, but I aim to further motivate the so-called Camp Doctrine by demonstrating its facility in application.

II On Perceptual Epistemology

For years, Susanna Siegel has defended a relatively unbounded view of perception whereby various cognitive processes can directly influence the phenomenology of an agent's perceptual experience. In her most recent book, Siegel devises a complex explanatory apparatus to provide much-needed heuristics for assessing the epistemological consequences of such cognitive influences on perceptual processes. In this section, I outline the innovative position Siegel dubs "perceptual conservatism" and contrast it with the more familiar model known as perceptual systems theory.

To begin, consider a standard example of perception:

MUG: *Rosalyn has forgotten that she left her mug of coffee on her desk. She looks around the room and notices her mug on the edge of her desk. As a result,*

(a) Rosalyn's perceptual experience includes her mug of coffee on her desk, and

(b) Rosalyn forms the occurrent belief "the mug is on the desk."⁷

The process involved in MUG that results in both (a) and (b) sees Rosalyn rapidly progressing through several distinct stages:

1. Rosalyn moves her eyes (and possibly her head) to scan the room in search of her coffee cup.
2. Her eyes fixate on the cup of coffee
3. Rosalyn [X-es] the cup.

⁷ Whether Rosalyn "forms" or "reforms" said belief is an irrelevant distinction for the present focus.

4. She forms the perceptual belief “the cup is on the desk.”

By “X-es” I mean only to indicate that a logical step exists between Rosalyn’s foveation on the cup and her formation of a perceptual belief about the cup; as will be explained, the exact nature of what it means to *X* is a key difference between perceptual systems theory and Siegel’s perceptual conservatism. Furthermore, although models of perception will rarely attribute much significance to the role of (1) (beyond treating it as one of many necessary physiological precursors to precipitate some instance of perception), such models can also be distinguished by their explanation of how the perceptual belief formed in (4) is justified and by which steps a model considers proper elements of “perceptual experience.” The remainder of this Section is devoted to explaining how PST and perceptual conservatism diverge on each of these explanations.

1. Perceptual Systems Theory and MUG

Consider an explanation of MUG from the perspective of Perceptual Systems Theory, which holds that a belief is a perceptual belief if and only if it is the output of a perceptual system (Lyons 2009, 98), where such a system is defined as:

a cognitive system that starts with the stimulation of sense organs by physical energy as input and processes information about the current environment, where none of the inputs to any of the subsystems are under the direct voluntary control of the larger organism (Lyons 2009, 94–98).⁸

The information-processing abilities of such systems produce two very different outputs: a sensory experience with phenomenological content (hereafter referred to more simply as a *sensation*) and a nonexperiential *percept* with conceptual content; these two outputs together

⁸ As will be explained more clearly in Section III, this characterization of PST follows a weakened version of Fodor’s modularity thesis (1983).

comprise the sum total of the agent's perceptual experience (Lyons 2015b: 156). So, this model would complete (3) in two steps:

- 3a. Rosalyn's visual perceptual system receives and discriminates light reflecting from the cup of coffee and its environmental surroundings as input.
- 3b. Rosalyn's visual perceptual system processes the information in (3a) and outputs:
 - i. A *sensation* (or collection of sensations) capturing the raw, qualitatively rich phenomenological features of Rosalyn's cup, such as its color, shading, and shape.
 - ii. A *percept* with the content "the mug is on the desk."

Notably, it is nomologically possible for the outputs of Rosalyn's perceptual system to come apart: she could have the sensation without the percept or receive the percept without the sensation. The former case (of having sensations without percepts) would be roughly what happens with unfamiliar perceptions, such as when an infant sees a coffee mug for the first time; the infant would experience the same qualitative features of the external world (the same colors and shapes), but would not be able to categorize those features as "coffee mug" because she has never before seen such an object. The latter case (of having percepts without sensations) is roughly what occurs in patients with blindsight who are able to reliably navigate their environment despite reporting an utter lack of visual sensations regarding their surroundings (Weiskrantz 1986).⁹

⁹ I say "roughly" because the reports of blindsighted patients typically indicate considerable hesitation and do not suggest a robust ability to perceive fine-grained details of their surroundings. More technically, the level of accuracy required to countenance the full separation of Rosalyn's percepts and sensations in (3b) would follow Block's imaginary condition of *superblindsight* wherein a patient themselves "contrasts what it is like to know visually about an X in his blind field and an X in his sighted field. There is something it is like to experience the latter, but not the former," they say (1995:233). Though conceivable, it is far from clear that superblindsight is nomologically possible.

Furthermore, Perceptual Systems Theory can recognize a reliabilist justification for perceptual beliefs formed in virtue of reliably-operating perceptual systems,¹⁰ leading to an additional clarification to the explanation of Roslyn’s perception of her coffee mug:

5. In virtue of the percept gained in (3b), the perceptual belief Rosalyn forms in (4) is epistemically justified.

Because her belief is directly based on the nondoxastic output of her perceptual systems, Rosalyn’s belief that “the mug is on the desk” is not based on any prior beliefs — that is to say, her belief about the mug is basic. Of course, Rosalyn’s justification for such a belief is subject to any number of potential defeaters, but, in the absence of such influences, the justified status of the belief remains.

2. Perceptual Conservatism and MUG

In contrast to this view, Siegel’s “perceptual conservatism” makes no clear distinction between sensations and percepts, focusing instead on the phenomenological category of “perceptual experience” as the purported ground for perceptual justification. A host of issues, ranging from longstanding questions about illusions and hallucinations to the relatively recent problem of cognitive penetration, precipitate Siegel’s need to construct a complicated explanatory framework to map how perceptual beliefs can ever be justified.

To begin, Siegel distinguishes perceptual *experience* from perceptual *judgment*, treating the former as “the conscious part of perception that [a subject] responds to when he forms his judgement” while the latter is what I have previously referred to as the process of perceptual belief–formation.¹¹ Consequently, perceptual experiences are phenomenologically rich and informationally saturated, but not in a manner that can be cleanly demarcated into sensations and

¹⁰ See also Goldman (1967).

¹¹ Specifically, she says “Perceptual judgment is a form of belief in which perceivers respond to the way things appear” (2017: 4).

percepts: instead, Siegel describes an individual's *overall* perceptual experience as "a collection of phenomenally conscious states," where *phenomenal states* are mental states "individuated by what it's like to be in them" and from which we learn about our surroundings (2011: 20). In a case like MUG, this view would consider Rosalyn's overall perceptual state as her full experience of her room as she's scanning it (including perceptual information from all of Rosalyn's senses); her *visual perceptual experience* is defined by the subset of her overall collection of phenomenal states comprised solely of visual phenomenal states.

Put differently, this view treats Rosalyn's experience in (1) as a messy collection of phenomenal states (some visual, some otherwise), while (2) defines a focusing of her perceptual experience onto the subset of those phenomenal states produced by the mug. Therefore, this model could treat [X-ing] in (3) as so:

3. Rosalyn has a visual perceptual experience of the mug.

From which, she then forms a perceptual belief in (4) — the process Siegel calls perceptual judgment. Altogether, Siegel's view treats Rosalyn's perceptual experience in (3) roughly as an form of what Sellars called the Given, insofar as the experience's non-epistemic phenomenal character simply provides a collection of facts about the mug to Rosalyn which thereby serve as a foundation from which she can form justified beliefs ([1956] 1997: 16).

Of course, Sellars' motivations for defining the Given were explicitly to debunk it for its epistemological weaknesses. Briefly, Sellars identifies a conflict between the claim that "sensations *detect* extramental particulars" and that "sensations *provide justification for beliefs about* extramental particulars." For sensations to do the former, sense experiences must contain a kind of non-epistemic content which can provide information about the world; for sensations to do the latter, their content must be epistemically empowered so as to be able to confer

justification to beliefs based on that content. According to Sellars, perceptual justification (so construed) leads to an infinite regress, given that the purported ability for sense experiences to confer justification requires the perceiving subject to already have “knowledge by acquaintance” of the object being perceived ([1956] 1997: 18). The canonical summary of this problem comes from Bonjour, who frames “epistemic content” as a matter of cognition:

...the givenist is caught in a fundamental dilemma: if his intuitions or immediate apprehensions are construed as cognitive, then they will be both capable of receiving justification and in need of it themselves; if they are non-cognitive, then they do not need justification but are also apparently incapable of providing it (1978: 11).¹²

One way of viewing Perceptual Systems Theory is as a model that cashes out the non-epistemic horn of this dilemma as a matter of percepts; Siegel’s perceptual conservatism instead grabs the other horn and argues that perceptual experiences can, in themselves, be epistemically evaluable and are, therefore, able to confer epistemic justification.¹³

This is the core notion of Siegel’s Rationality of Perception (RP) thesis:

RP: Both perceptual experiences and the processes by which they arise can be rational or irrational (2017: 15).

Siegel defines x as potentially rational *iff* x is “epistemically appraisable” in a manner that could be “appropriate or inappropriate, relative to...norms of rationality” (15–16). Typically, this applies uniquely to various propositional attitudes and their transitional processes (such as an inference which leads one to believe a proposition); riffing off a line from Donald Davidson,

¹² For a slightly different summary of the so-called Sellarsian dilemma that maintains the key focus on belief-justification, see (Lyons 2016b: 1062).

¹³ Siegel says as much almost explicitly, explaining, “In the last decades of the twentieth century, many analytic philosophers rejected the second horn of Sellars’s dilemma, on the grounds that experiences could both have accuracy conditions and provide justification for subsequent beliefs, without thereby needing to be justified by anything else. I think this response succeeds. But all forms of the response I know of stopped short of claiming that experiences with these features could themselves manifest an epistemic status. The idea that experiences can be epistemically charged sits comfortably with the prominent response to Sellars’s dilemma, and takes it one step further” (2017: 29).

Siegel explains that the RP thesis intentionally “expands the house of reason” in that it “extends the domain of rationality to transitions and states that were previously assumed to be entirely a-rational” (19).¹⁴ Namely, it treats phenomenologically rich perceptual experiences (and their related transitional processes) as epistemically appraisable.

However, Siegel is clear that perceptual experiences are not measurable in the same way as beliefs or other attitudes; whereas *justification* is the epistemic metric used to assess the contribution of a belief to an agent’s overall rational standing, Siegel posits a different, distinct heuristic for the epistemic status of perceptual states: epistemic charge.¹⁵ Patterned along a metaphor with electric charge, epistemic charge is a valenced concept by which every perceptual experience can be measured positively or negative in comparison to a neutral baseline (27). In cases where a perceptual experience is at the epistemic baseline, it possess sufficient *epistemic power* to support well-founded beliefs; epistemic *downgrade* occurs when a perceptual experience loses its baseline epistemic power and, thereby, is less able to support well-founded beliefs (65); epistemic *upgrade* occurs when a perceptual experience is *enriched* beyond the epistemic baseline, thereby increasing the justification of beliefs formed on the basis of that experience (75–77).

Consider two cases as illustration:

MOE: Moe is a human being standing in front of a normal, round soccer ball. When Moe looks at the soccer ball, he forms the belief (on the basis of how the ball looks to him) that the ball is round.

SHMOE: Shmoe is Moe’s identical twin, separated from Moe at birth and raised in extremely unusual conditions such that he believes that

¹⁴ For the context of the original “house of reasons” example, see Davidson (2004: 169).

¹⁵ Admittedly, Siegel explicitly remains neutral on just how distinct epistemic charge is supposed to be from justification (28–29) and, at least once, offers a justified belief as an example of one kind of epistemically charged mental state (25). Given that her main point is simply to establish that epistemic charge applies to non-doxastic mental states, I’m not sure this ambiguity is ultimately significant for her argument (however exegetically noteworthy it might be).

soccer balls are not round. When Shmoe stands next to Moe and looks at the same normal, round soccer ball, Shmoe forms the belief (on the basis of how the ball looks to him *as affected by the contents of his prior outlook*) that the ball is not round.

In the absence of defeaters or other unusual situations that could affect Moe's visual experience, Siegel argues that Moe's visual experience of the ball possess baseline epistemic power and, therefore, allows Moe to form a well-formed belief that the ball is round; furthermore, insofar as Moe's experience is epistemically empowered, Moe is thereby capable of making all manner of perceptual judgments on the basis of that experience. However, because Shmoe's prior outlook infects his visual experience and causes him to view the ball differently, his visual experience is downgraded below the baseline and thereby precludes Shmoe from making justified perceptual judgments. According to Siegel's model, both Moe and Shmoe have formed well-founded beliefs (insofar as they are *prima facie* justified in believing what they see), but the negatively-charging effect of Shmoe's prior outlook on his perceptual experience defeats his ability to form proper perceptual beliefs.

Importantly, epistemic charge is not simply identical with the justification of beliefs, but is a *sui generis* property by which a perceptual experience is supposed to be, in itself, epistemically appraisable.¹⁶ Furthermore, Siegel argues that the presentational phenomenal character of perceptual experience that "purports to characterize how things in the external world are" is what grounds the epistemic charge of a perceptual experience (45). This means that, because all perceptual experiences have presentational phenomenal character, all experiences have an epistemic charge.¹⁷ And, so long as it is *simply* the presentational phenomenal character

¹⁶ Per Siegel, experiences with a charged status "are not merely an enabling condition for other mental states to manifest such a status. Nor are they merely contributors to determining the epistemic status of beliefs as either well or poorly justified, though they play this role as well. In addition to this role, they manifest an epistemic standing in themselves" (26–27).

¹⁷ This is the conclusion of Siegel's "Phenomenal Ground" Argument (45).

modulating the epistemic charge of an experience, the experience's charge remains at or above the baseline; as Siegel explains, "in the pervasive situations in which it is reasonable to believe your eyes, if perceptual experiences have any epistemic charge at all, they have a positive epistemic charge" (46).

Therefore, just as justification makes a contribution to determining an agent's overall rational standing (by, for example, judging that someone who believes p without justification is "irrational"), so too does epistemic charge. Siegel gives the following example of how this might play out:

Consider two equally rational subjects, call them S^- and S^+ , where S^- has an ill-founded belief that p and S^+ has a well-founded belief that p . The two subjects could end up equally rational, because [the epistemic charge of] S^- 's other mental states compensate for the illfoundedness of her belief that p . So we can't read off from the fact that S^- has an ill-founded belief that p that she is less rational than a subject with a well-founded belief that p (42).

Ultimately, epistemic charge operates by making a *pro tanto* contribution to a person's overall rational standing.

On this view, because all experiences have *some* epistemic charge, the etiology of a given experience — such as whether it has been influenced by inference or some other feature of a subject's doxastic system — becomes crucial for determining the epistemic power it has license to confer on beliefs (44). Although, as discussed above, most perceptual experiences unperturbed by unusual conditions or defeaters will give a subject good grounds for forming well-founded beliefs about the world, if "defeaters or epistemically bad-making etiologies" influence how a subject comes to have a perceptual experience, then "[e]ither kind of bad-making factor would

prevent the phenomenal character of experience from bestowing baseline positive charge” (47).¹⁸ If I see a red statue, then my visual perceptual experience will carry the epistemic power to justify beliefs like “That statue is red,” but if it is the case that a red light is shining on the statue (thereby only making it look red to me), then this fact will disempower my perceptual experience, regardless of whether or not I am aware of the red light’s effect on the scene — put differently, my experience will be negatively charged. Also, recall how the influence of Shmoe’s prior outlook on his visual experience similarly undercut the ability of his visual phenomenology to license justified beliefs about the roundness of the soccer ball.

Altogether, epistemic downgrade — whereby a subject’s perceptual experience carries negative epistemic charge — is a way of identifying the epistemic perniciousness of cognitive penetration, when the phenomenal character of a token perceptual experience is affected by one or more preceding psychological states, such as beliefs, emotions, or desires. Siegel concludes that, in virtue of tracking the epistemic charge of experiential states, epistemologists can measure the effect of one’s “prior outlook” on the contents of one’s perceptual experiences rather than simply rejecting foundationalism (49–50).

What this then amounts to is that a perceptual experience formed at least in part on the basis of prior beliefs — and not on purely synchronic sensory inputs — is subject to what Siegel dubs “perceptual hijacking” (5). Ultimately, Siegel’s rational perceiver walks a fine line: “Perception goes well, either as experience or judgment, when perceptual inputs are given proper weight. And perception goes badly, when perceptual judgment or experience is hijacked by one’s prior outlook” (5).

¹⁸ Although when introducing the concepts, Siegel treats “the epistemic baseline” as something distinct from “positive epistemic charge,” she tends to later conflate the two, treating “positive baseline” as a functionally similar to “empowered to confer epistemic justification.”

3. Perceptual Systems Theory, Perceptual Conservatism, PINE, and SCOPE

Consider again the PINE case. Perceptual systems theory could argue that, over time, Calvin either learns to more quickly cognize the percepts output by his visual perceptual system or, on a perceptual learning thesis, trains the visual perceptual systems themselves to output richer percepts related to pine trees in a manner such that their uptake into conceptual processing is faster. Either way, such a training process is effortful and time-consuming, at least in part because it requires the enrichment of <pine tree> available to Calvin. Notably, although Calvin *might* report a phenomenal contrast in his visual experience of the forest (or, more specifically, of pine trees) from the beginning of his project to the end, he also might *not* — such a shift, if it occurs, would be a merely incidental byproduct of the process of perceptual learning.¹⁹ Provided that his perceptual systems are nonetheless trained in a manner that allows him to reliably form the right kinds of beliefs about pine trees (insofar as his systems are outputting pine-tree-related-percepts in the right kind of way), then the experiential or phenomenological character of his perceptual experiences is explanatorily superfluous.²⁰

In contrast, Siegel's perceptual conservatism could say that PINE evidences an enrichment of Calvin's overall outlook on the world such that he gains <pine tree> in a manner which comes to affect his visual processes (perhaps through the habitual tokening of beliefs like "there are pine trees in this forest"). Insofar as Calvin comes to hold certain beliefs about the visual appearance of pine trees, those features of his outlook could enrich his perceptual experience, predisposing him to better notice the pine trees that really are in his environment. The influence of his doxastic states on his perceptual experience modulates the experience's

¹⁹ Here, I have in mind a retort from Fodor (1988) directed at Churchland about diachronic penetration that will be discussed in Section III.

²⁰ Put differently, I'm arguing that philosophical zombies could undergo the process of perceptual learning, despite their complete lack of what Siegel calls "presentational phenomenal character."

charge, but because the accuracy conditions of the content of his experience continue to track his environment, Calvin's perceptual experience is *upgraded* — beliefs based on this visual experience are well-founded and Calvin's overall rational standing is uncorrupted.²¹

Because the preformationist thesis about biological development is false, the SCOPE case cannot work in the same manner as PINE. Because there is no embryo to even potentially perceive inside a sperm cell, PST entails that a perceptual system will output no embryo-belief-supporting percepts; Thomas' perceptual experience is either illusory (whereby he misinterprets the character of the mental representation formed on the basis of the actual percepts generated by his visual perceptual system) or a hallucination (whereby he forms an unjustified mental representation, despite the absence of output percepts). Either way, beliefs formed by Thomas about <the embryo on his microscope slide> as a consequence of his perceptual experience will ultimately be unjustified, so his overall rational standing will be undermined. Given the etiology of how he might report his perceptual experience, this would be a textbook case of perception's cognitive penetration brought about by wishful thinking (sometimes referred to as "wishful seeing"); although Thomas might form all manner of beliefs (perceptual and otherwise) about the contents of his slide, the actual content of his perceptions (as determined by the percepts generated by his perceptual systems) is different.²²

In contrast, perceptual conservatism explains SCOPE as a case where Thomas' perceptual experience is epistemically *downgraded*, due to the influence of his prior outlook *vis a vis* his commitment to preformationism. On this view, Thomas' report that he has a perceptual

²¹ Siegel(2017) offers roughly this example, saying "For instance, consider a tree expert and a non-expert who look at the same pine tree, and the expert gets more justification to believe it's a pine tree than the non-expert gets, because the expert's experience has the content *x is a pine tree*" (75).

²² This is to be contrasted with Siegel's (2017b) notion of wishful seeing, whereby an agent's phenomenological experience is *genuinely different* as a result of their doxastic influence; as she says, "Some cases of wishful seeing are cases of wishfully seeing that P (where that state is factive), others are cases of wishfully experiencing that P" (409).

experience as-of an embryo inside of a sperm cell is taken for granted and treated as a genuine case of perception, but because the epistemic profile of the experience is below the baseline (due to the corrupting influence of his prior outlook that does not accurately capture the nature of the world), any beliefs Thomas forms on the basis of this perceptual experience are *not* well-founded. Despite the fact that Thomas not only genuinely believes that he sees an embryo where there is, objectively, no embryo, but (on this view) Thomas *genuinely sees an embryo* (where “see” means “has a visual perceptual experience as of”), the downgraded nature of that experience would defeat the epistemic power of the experience’s phenomenal character and, thereby, undermine Thomas’ overall rational standing.

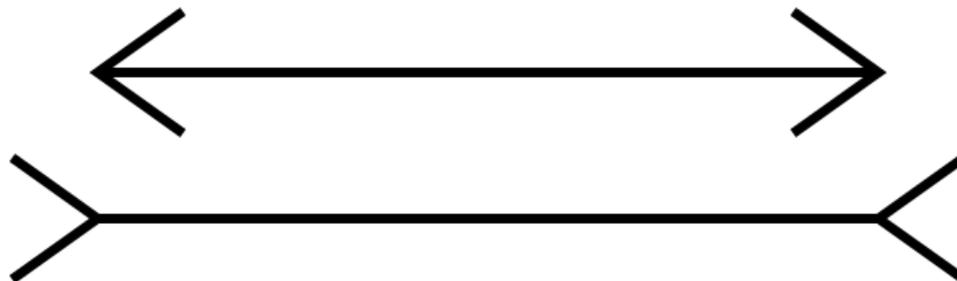
Early on in *Empiricism and the Philosophy of Mind*, Sellars frames his project around a concern for distinguishing between a subject “sensing sense contents” and a subject “being conscious” ([1956] 1997: 20); he treats a conflation of one’s sensory experience of phenomenal states with one’s cognitive judgments about those states as something akin to the naturalistic fallacy in ethics — simply treating experiences as epistemic objects (of the sort that agents cognize as a function of “being conscious”) is to treat experiences as something that they fundamentally are not. Whereas perceptual systems theory can respond to the Myth of the Given by maintaining Sellars’ distinction (by disambiguating percepts from perceptual experiences), perceptual conservatism’s emphasis on experiential phenomenal character offers no such option and argues instead that, contra Sellars (and many others in the history of Western philosophy), experiences *are* epistemic objects. To then explain how an agent’s overall rational standing might be preserved even in a case of cognitive penetration (or otherwise “hijacked” experiences), Siegel has constructed the complex conceptual apparatus described in this section, positing the property of epistemic charge, the modulating effect of inference on an experience’s epistemic

power, and the significance of an experience's etiology for determining the well-foundedness of beliefs based on its phenomenal character.

Having introduced a basic summary of the view in this section, I will now turn to an appraisal of perceptual conservatism as a model in Section III; in particular, I will consider a problem case that perceptual conservatism seems unable to sufficiently explain.

III Why Perception is Nonrational

The previous section introduced Siegel's model of perceptual conservatism, motivated by the rationality of perception thesis, and contrasted it with an alternative model of perceptual epistemology — perceptual system theory. In this section, I will argue that the recalcitrance of perceptual illusions is evidence that the rationality of perception thesis is false, thereby undermining perceptual conservatism and privileging perceptual systems theory.



*Fig. 1: The Müller-Lyer Illusion*²³

Consider the Müller-Lyer illusion (MLI), where two lines of identical length appear to be different lengths because of the differently angled wings attached to their endpoints. Regardless of how familiar a subject is with the MLI, the experience of the illusory effect persists; that is, the effect of the MLI is recalcitrant, no matter what kind or how many beliefs a subject might

²³ As created by Franz Carl Müller-Lyer; see https://en.wikipedia.org/wiki/Müller-Lyer_illusion.

have about the psychological mechanisms relevant to the phenomenal experience of the lines and regardless of how confidently the subject believes that, in fact, the lines are not different lengths.

One explanation of the obdurate phenomenology of the MLI relies on the distinction between percepts and concepts, already discussed in Section II. Consider the following case:

MAX: When Max — a person who has never before heard of nor seen the MLI — looks at the MLI for the first time, he lacks <MLI>, but his perceptual systems nevertheless output percepts *of* the MLI that trigger and give him *prima facie* justification for beliefs like the following:

- 1) “Those two lines *seem* to me to be of different lengths.”
- 2) “Those two lines *are* of different lengths.”

To use language introduced in Section II, (2) might be an example of a perceptual belief (or, in Siegel’s terminology, a perceptual judgment); it is not immediately clear, though, what the *seeming* described by (1) is.

Causal theories of perception would say that the MLI (or Max’s percepts of the MLI) reliably cause (1), so Max’s belief in (1) is justified (Goldman 1967; Grice and White 1961);²⁴ once Max learns the nature of the MLI, he will then have defeaters that block the (unnoticeable) inference from (1) to (2), but such defeaters do not affect his justification for (1). Following Turri (2011), if Max’s percepts “non-deviantly cause” Max to believe (1), then we can say that Max always has a reason to believe (1); insofar as Max is incapable of directly controlling the outputs of his perceptual systems, he is likewise incapable of controlling what sorts of beliefs those outputs invariably trigger and so, if he is operating rationally, will intractably believe (1) and, thereby, make the inference to (2).²⁵

²⁴ More recently, Turri (2019) has also argued for experimental support of such theories.

²⁵ Naturally, Max could move his head or willfully fix his gaze elsewhere, but this would result in an *indirect* alteration of the outputs of his perceptual systems.

However, if at least some of a perceptual system's outputs are beliefs, then the story is even less contingent (and, therefore, less open to potential defeaters); the output of Max's perceptual systems could include (1) directly, so Max will intractably believe (1), provided that his perceptual processes are operating in their reliable fashion.²⁶ As Glüer (2018: 3000) explains, Max can easily believe (1) while believing that (2) is false by recognizing that (1) is merely making a phenomenological claim that does not conflict with his background beliefs that deny (2). So, this model treats a seeming of (1) as the perceptual belief and (2) as a belief inferred from (1), but subject to potential defeaters. Importantly, though, while denying the link from (1) to (2) is under Max's control, his belief in (1) is not: Max's rational processes have nothing to do with his process of perception.²⁷ Indeed, this kind of view effectively evacuates the search for epistemic reasons for perceptual beliefs altogether, instead placing the weight of perceptual justification squarely on the reliability of the perceptual process alone.^{28,29}

Either way, at this point, the most important thing is that, regardless of why Max's belief in (1) is obdurate and automatic, Max's belief in (2) is *neither*: it is, in one way or another, inferentially based on his belief in (1) and is, therefore distinct from it.³⁰ Upon learning that, in fact, belief (2) is false, Max is not required to likewise disavow (1) — indeed, for the reasons described above, he cannot.

²⁶ For example, see the discussion in Glüer (2009) about a “phenomenal doxastic” account of perceptual experience; I take it that the distinction I draw here between (1) and (2) is roughly akin to such an account.

²⁷ Lyons (2009: 72) goes a step further to argue that some percepts could even be token-identical with beliefs like (2), further tightening the grip of justification for perceptual beliefs.

²⁸ For more on this, see Chapter Nine of Kornblith (2015), particularly pages 199–200.

²⁹ To be clear, Glüer's intentionalist explanation is not a reliabilist one, instead arguing that experiences can defeasibly justify beliefs and explaining illusions (and other odd experiences) by distinguishing between “phenomenal” and “standard” semantics for seemings and beliefs, respectively (2018: 2998).

³⁰ On inferential basing relationships (as opposed to non-inferential basing relations), see Korcz (2019). In short, Korcz argues that basing relations are sometimes, but not always, inferences and that inferences sometimes, but not always, depend on basing relations, but “every inferentially justified belief involves an inference coinciding with a basing relation” (120).

1. Perceptual Systems as Modular Input Systems

According to Perceptual Systems Theory, the intractability of perceptual beliefs like (1) stems from the fact that perceptual systems are what Fodor (1983) calls modular “input systems.” Although he never offers a specific definition of what a ‘module’ is, Fodor does give several examples and discuss various common features of modules while saying that “[r]oughly, modular cognitive systems are domain specific, innately specified, hardwired, autonomous, and not assembled” and, most importantly, they function under conditions of “informational encapsulation” (1983: 37).³¹ As Lyons (2015a: 103–104) points out, Fodor stresses that these diagnostic features of modules are not categorical, necessary, or sufficient conditions for identifying modules as such; it is merely the case that modular systems exhibit these features “to some interesting extent.” The visual systems relevant for perceiving the MLI certainly fall under Fodor’s definition to some interesting extent: they pertain uniquely to the visual domain, they are innate and hardwired by an agent’s biology (including a particular neural architecture), they operate automatically without doxastic initiation and also, crucially, without doxastic influence: in short, I do not need to think about what’s in front of me in order to see it, I simply need to look.

In fact, Fodor (1984) deploys precisely this sort of framework for explaining the phenomenology of the MLI. At the time, he was engaged in a lively debate about the purported theory-neutrality of observation; arguing particularly against critics of theory-neutrality like Churchland (1979), Fodor points to paradigmatic examples of “curious and persuasive perceptual *implasticities*” as proof that a subject’s perceptual systems are informationally encapsulated —

³¹ Slightly more technically, input systems perform domain-specific, informationally encapsulated operations that are mandatory, rapid, and often unconscious (but which follow predictable breakdown patterns, paces, and sequences), these computations are associated with fixed elements of a subject’s neural architecture and inevitably produce only shallow outputs (Fodor 1983: 47–101).

the MLI is one such example (1984: 34). To be clear: the Fodor-Churchland debate about ‘observation’ (in the technical sense of the term as they use it) extends beyond questions of simple perception as discussed here, but Fodor’s concern to disambiguate “belief fixation directly consequent upon the activation of the senses” and “belief fixation via inference from beliefs previously held” is precisely relevant to the present discussion (1984: 23): the former concerns the process by which Max forms a seeming like (1), while the latter characterizes Max’s belief in (2).

For Fodor, the exact sequence of a perceptual system’s operation or the “process wherein an organism assigns probable distal causes to the proximal stimulations it encounters”³² might well rely on the inferential, problem-solving processes of Churchland-style plastic architecture, but the fact that a subject is utterly ignorant of the processes internal to her perceptual system and the fact that those processes are not affected by “ALL (or, anyhow, arbitrarily much)” of her other, higher-order cognitive states is sufficient to consider them encapsulated in Fodor’s sense (1984: 30–35). Even if the details of the inner workings of Max’s perceptual systems that output (1) follow various inference-like processes — á la the 2½D-sketch-building model proposed by Marr (1982) or the early/late vision distinction defended by Pylyshyn (1999) — their overall resistance to Max’s conscious operation — such as his willful rejection of (2) — preserves their status as properly encapsulated “to some interesting extent.”³³

³² By “proximal stimulations,” Fodor means something like the activation of sensory transducers.

³³ Fodor (1983: 78) makes roughly the same point by admitting that “...mechanisms *internal* to a module might contrive to, as it were, mimic the effects of cognitive penetration” without actually amounting to real cognitive penetration.

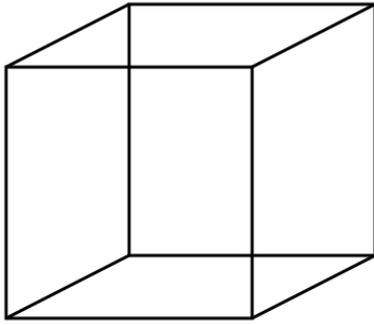


Fig. 2: *The Necker Cube*³⁴

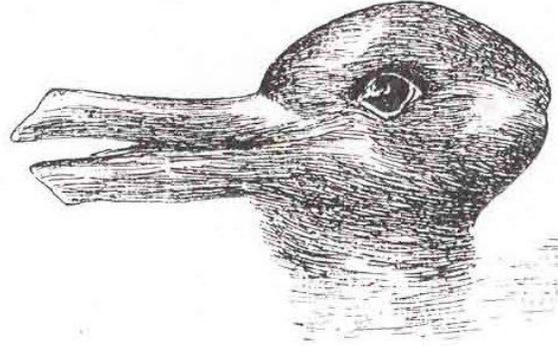


Fig. 3: *The Duck-Rabbit*³⁵

In a response to Fodor, Churchland capitalizes on Fodor's concession of perception's inferential components to push back on the encapsulation thesis, but then admits that an experience of the MLI cannot be willfully altered in a moment by a subject (as an experience of an ambiguous figure like the Necker cube [Fig. 2] or the duck-rabbit head [Fig. 3] can); Churchland suggests that this is because the phenomenology of the MLI is "an incidental consequence of a long period of perceptual training on certain typical kinds of perceptual problems," such that only "a long period of training in an environment of a quite different perceptual character would produce a subject free from that particular illusion" (1988: 174). Fodor's response to Churchland points out that such a "long period of training" would demonstrate *diachronic*, not *synchronic*, penetration. In part, the issue of the difference is a matter of definition ("how much diachronic penetration would pose the same epistemological problems as synchronic penetration?" for example), but Fodor also indicates that the evidence for synchronic penetration is problematically unclear, saying (in discussing the purported "local" effects of beliefs on perception), "[t]hese might not even be perceptual effects of acquiring

³⁴ Named for Louis Albert Necker; see https://en.wikipedia.org/wiki/Necker_cube.

³⁵ A common example, available at https://commons.wikimedia.org/wiki/File:Duck-Rabbit_illusion.jpg.

beliefs; perhaps they're perceptual effects of *having the experiences* in virtue of which the beliefs are acquired” (1988: 192).

Two things are important about this final comment from Fodor: firstly, his discussion of diachronic penetration appears functionally akin to the process of perceptual learning already mentioned in Section I: both entail the slow, laborious change in a perceptual system’s output driven by a change in a subject’s environment. Churchland treats ‘diachronic penetration of perception’ and ‘perceptual learning’ as effectively synonymous, explaining that “[o]ur perceptual, practical, and social environment shapes our perceptual capacities mightily, especially in their early stages of development, and this suggests that different courses of learning would produce interestingly different perceptual capacities” (1988: 176). This means, as explained in Section I, diachronic penetration (i.e. perceptual learning) is distinctly *not* synchronic cognitive penetration and Churchland’s conclusion that, because “diachronic penetration is commonplace” that cognitive penetration is commonplace and Fodor’s modularity thesis is undermined (177), is simply incorrect.

But, importantly, Fodor’s treatment of perceptual learning is *also* problematic. Because his definition of modularity commits him to affirming that modular systems have “fixed neural architectures,” Fodor felt compelled to concede that “the involvement of certain sorts of feedback in the operation of input systems would be incompatible with their modularity” (1983: 66). That is, if it could be demonstrated that modular systems could *change*, then Fodor would be apt to take that as evidence *against* the modularity thesis (or, at the very least, the encapsulation thesis); as Raftopolous puts it, “[w]ell, the findings are out and suggest that perceptual systems are diachronically, meaning in the long run, open to some rewiring of the patterns of their neural connectivity, as a result of learning. In other words, these systems are to some extent plastic”

(2001: 443). However, Raftopolous goes on to explain precisely why this does *not* entail that said systems are cognitively penetrable: his review of the scientific literature indicates precisely that such neural rewiring is reliably provokable only by so-called “slow learning” processes that are grounded in the alteration of “low-level, stimulus-dependent visual processing stages,” in a manner that is “independent of cortico-limbic processing, which is responsible for top-down processes,” and is “independent of factors involving semantic associations” (443–444). In short, Raftopoulos indicates that the neural architecture of perceptual systems is *not* innately fixed, but is sensitive to the complicated, gradual process of perceptual learning; “slow learning takes place under specific retinal input and task-dependent conditions. These factors do not involve cognitive influences, and therefore, do not imply the cognitive penetrability of perception” (444).

So, when Fodor (1983: 33) admits that “children, having had less experience with edges and corners than adults, are correspondingly less susceptible to the [MLI],” this need not undermine the overall analysis of Max’s perception of the MLI — and his unavoidable fixation of belief in (1) — as a modularly automatic process. Similarly, evidence that population-level susceptibility to perceptual illusions like the MLI is variable cross-culturally offer further evidence along these lines (Heinrich, Heine, and Norenzayan 2010); such populations demonstrate Churchland’s point that different ecological conditions would provoke a different contingent process of diachronic penetration such that the experience of the MLI is different (1988: 176). But, crucially, *neither* case evidences *cognitive penetrability* as the culprit for the difference in either phenomenological experience or the simple fixation of belief — in both cases, it is *perceptual learning* that is responsible.

To indicate that they are subject only to what he calls the “off-line penetrability” of slow-learning processes, Raftopolous labels perceptual systems *semien*capsulated (445), thereby

suggesting that they carry many — but not all — of the diagnostic features of modularity which Fodor insisted input systems need to display “to some interesting extent.” Such language dovetails with the more fine-grained treatment of perceptual systems found in Lyons (2015a) that distinguishes systemic *isolability* (with respect to task-computation) from systemic *encapsulation*: the former “is about system boundaries, about what is required to have an intact computational device” while the latter is “a matter of where a given system gets its inputs from” (106).

On Lyons’ framing, perceptual modules might not be merely semiencapsulated, but entirely *unencapsulated* and, nevertheless, function to reliably produce true perceptual beliefs. Importantly, this could apply, in the case of Max’s experience of the MLI, to perceptual beliefs in both the sense of (1) *and* (2) because, *pace* Lyons, the epistemological problems of cognitive penetration on the fixation of perceptual beliefs is ultimately no different than the epistemological problem of the fixation of *many* beliefs (2015a: 118). Rather than encapsulation, the truly worrisome problem is the *introspective opacity* of perceptual systems (an opacity matched by many belief-fixing systems, such as intuition) given that it limits our awareness of a produced belief’s justificatory status (118). Insofar as cognitive penetration can *promote* the overall reliability of perceptual processes, then Lyons (2011) is happy to recognize its positive contribution to perception, but only in virtue of the fact that such cognitive penetration *promotes true belief-fixation* — the alteration of a subject’s visual perceptual experience is entirely beside the point.

Which returns us to the first important thing about what Fodor said regarding synchronic cognitive penetration: recall that he *also* makes roughly the same point as Lyons, arguing that the effects of synchronic penetration “might not even be perceptual effects of acquiring *beliefs*;

perhaps they're perceptual effects of *having the experiences* in virtue of which the beliefs are acquired” (1988: 192). By this, I take Fodor to mean that having phenomenological experiences *is not necessary* for proper belief-fixation to obtain; he is driving the same sort of wedge between belief-fixation *in toto* and phenomenological experience *simpliciter* already described as a Sellarsian-style Myth of the Given.³⁶ As a nascent perceptual systems theorist, Fodor could respond to the Sellarsian dilemma in roughly the way outlined in Section II, but the distinction between visual phenomenology and the beliefs a subject *bases* on that phenomenology is still important to maintain. If “putting us in touch with the world is the point of perception *epistemologically speaking,*”³⁷ then this is a matter of reliable belief-formation — any sort of qualitative phenomenology is simply an added bonus.

2. Perceptual Conservatism and Recalcitrant Illusions

The second important thing about Fodor’s distinction between “perceptual effects of acquiring beliefs” and the “perceptual effects of *having the experiences* in virtue of which the beliefs are acquired” follows from the first: if justifiable belief-fixation is indeed divorceable from qualitative experience, then phenomenology is not only not necessary for grounding justified beliefs, but — insofar as it is unreliable — it is not sufficient for grounding them either.

Consider how Siegel’s perceptual conservatism must explain Max’s experience of the MLI (when he looks at the MLI for the first time): by looking at the two lines, Max experiences a collection of phenomenal states which each have a given presentational phenomenal character, the sum of which composes a visual experience of the lines. This perceptual experience is both

³⁶ Even if Fodor meant “experience” in a less technical sense here, he nevertheless divorces the phenomenological component of a personal experience from the agent’s beliefs about both the experience and the experience’s phenomenology, thereby setting up later thinkers (such as Lyons) to expand this distinction more robustly (as described in this subsection).

³⁷ Lyons (2015a: 119)

phenomenologically rich and informationally saturated which, in the absence of “defeaters or epistemically bad-making etiologies” possesses sufficient positive epistemic charge to ground beliefs (2017: 47). Because Max is unfamiliar with the MLI, he has no reason to doubt his generally-reliable perceptual processes, so, on the basis of his visual phenomenology alone, Max feels empowered to conclude both:

- 1) “Those two lines *seem* to me to be of different lengths.”
- 2) “Those two lines *are* of different lengths.”

Of course, (2) is false, but Max’s ignorance of the MLI means that he will have no reason *not* to make the inference from (1) to (2). Only upon learning about the MLI will Max become aware of the relevant defeaters which undercut the inference from (1) to (2) in the first place. However, if phenomenology were sufficient to ground beliefs, then no further restrictions on visual phenomenal character (such as “in the absence of defeaters or epistemically bad-making etiologies”) would be needed. If only certain kinds of phenomenal states are sufficient for grounding belief, then phenomenology, full stop, is not sufficient.³⁸

Consider the subset of phenomenal states Siegel focuses on: normal cases without defeaters. Notably, such as in the case of MOE discussed above, the presentational phenomenal character of a visual perceptual experience offers direct justification for beliefs like (2) without needing the intermediate step of (1) or anything else; as Siegel explains, “the experience with content *p* could provide immediate justification for believing *p*” (48).³⁹ Calling such experiences both “self-justifying” and “self-charged” (48–49) Siegel again deploys phenomenal conservatism’s notion of epistemic charge to explain how phenomenal states contribute to an

³⁸ A point Siegel seems to make as well in suggesting that “experience, all by itself, is a bearer of epistemic power, but it is never by itself a bearer of baseline epistemic power” (2017: 64).

³⁹ This is also making use of what Siegel calls the “intermediate” role of phenomenal character (47).

agent's overall rational standing: in virtue of their phenomenology, they empower a perceiver to rightfully form perceptual beliefs.

So, on this view, the MLI is an unusual (and comparatively rare) aberration from how our perceptual processes typically work — a point Churchland made decades ago (1988: 174). Max's visual system is triggered by the two lines in the normal way that extramental objects trigger such systems such that he has a particular perceptual experience as-of the lines being of different lengths, but Siegel is clear that such perceptual experience is always distinct from perceptual *judgment*. Although perceptual experiences are not consciously alterable (Max could not, by some sheer force of will, change how the lines of the MLI appear to him), such judgment is “a form of belief in which perceivers respond to the way things appear” and this response *is* under Max's conscious control (2017: 4). Put differently, Siegel would say that, although Max cannot consciously alter his belief in (1), he can alter his belief in (2) — (1) constitutes Max's perceptual experience and (2) constitutes his perceptual judgment — and this disavowal ability is all one needs in order to protect his overall rational standing from such an unusual illusion. As Siegel explains, “If you know the lines are not the way they appear, and you cease to rely on the experience, then there is no further rational adjustment to be made in the situation” (2017: 36).⁴⁰

To Siegel, the more interesting (and problematic) cases are those of perceptual hijacking, where a subject's psychological mental states (as packaged within their “outlook” on the world) alter how things seem to a subject — as in the above case of SHMOE. By “outlook,” Siegel means the full collection of a subject's attitudes, beliefs, hypotheses, knowledge, desires, traits, and moods and any other relevantly similar mental states that constitute a subject's perspective on the world (1). In hijacked cases, a subject's outlook affects *either* that subject's perceptual experiences *or* their perceptual judgments such that their beliefs about the ways things *seem to*

⁴⁰ Notably, this is the only sentence in *The Rationality of Perception* that discusses the MLI.

be or the way things *are* in the world are different than they would be without the influence of such an outlook. So, Siegel further distinguishes between a subject's *considered* outlook — the collection of rationally assessed mental states that the subject consciously affirms — and their *complete* outlook, which includes all elements of one's considered outlook, plus all disavowed experiential states (37).

Consider a subject born into a virulently racist context whose outlook has been shaped in a racist way. Siegel argues that such a subject, when looking at another person, might not only make racist perceptual judgments, but might well have racist *perceptual experiences* as well. She gives an example of a defendant who attacked a man she perceived to be dangerous because the man appeared (to the defendant) to be holding a gun (even though he was not). If the defendant had a perceptual experience as-of the man holding a gun and therefore believed the man to be dangerous, and if her experience was infected by her racist outlook, then this is a case of perceptual hijacking: both the defendant's beliefs *and* the defendant's experience are irrational. Siegel's explanation is worth quoting at length:

If the reasonableness of the person depends only on the interface between experience and subsequent belief, and not on the psychological background of the experience, then the belief that the man is dangerous (because he is holding a gun) might seem to be reasonable. But if the gun-experience itself detracts from the subject's rational standing because it is inferred from an ill-founded presumption, then when we assess what a reasonable person under similar circumstances would believe, we need not hold constant their experience. A reasonable person in similar circumstances would not have an experience that they inferred from the ill-founded presumption. Just as a reasonable person's beliefs would be by and large shaped by reasonable presumptions, their perceptual experiences would be shaped that way as well (24).

Siegel's entire project (developing the concepts of epistemic charge, downgrade, upgrade, and the rest) is precisely to explain why such a racist perceptual experience is irrational: because it has been unduly influenced by the subject's outlook. So, if the subject with a racist outlook has a

racist perceptual experience, but then disavows (or avoids) racist beliefs based on that experience, their considered outlook might not be racist, but their *complete* outlook would remain infected with racism, nevertheless.

Siegel's intricate conceptual apparatus for measuring the contributions (positive and negative) of a perceptual experience to an agent's overall rational standing is precisely the sort of heuristic necessary to preserve "justified belief" about extramental objects *if* outlooks can affect perception as described. But, however rare and unusual of a case it might be, the reality of the MLI (and other seemingly-intractable illusions) offers a Gordian knot for this sort of outlook-penetration of perception that cannot be unraveled. To illustrate this, consider one last case:

SHMAX: Shmax is a person who is extremely familiar with the MLI, having viewed it many times, under many conditions, and has even studied it philosophically, thereby robustly developing his <MLI>. Nevertheless, when Shmax looks at the MLI, his visual systems trigger him to believe:

1) "Those two lines *seem* to me to be of different lengths."

But Shmax's prior outlook, which includes knowledge of the MLI, prevents him from believing:

2) "Those two lines *are* of different lengths."

Clearly, as described above, perceptual conservatism can explain why Shmax disavows (2) — his considered outlook includes the belief that the lines of the MLI are equal, even though his complete outlook also includes the disavowed perceptual experience of the lines being of different lengths.

However, a key problem for perceptual conservatism is in Shmax's perceptual *experience* as captured by (1), for it is at this point that Fodor's question from "Observation Reconsidered" looms:

The Muller-Lyre (sic) is a *familiar* illusion; the news has pretty well gotten around by now. So, it's part of the 'background theory' of anybody who lives in this culture and is at all into pop psychology that displays like [Fig. 1] are in fact misleading and that it always turns out, on measurement, that the center lines of the arrows are the same length. Query: *Why isn't perception penetrated by THAT piece of background theory?* Why, that is, doesn't *knowing* that the lines are the same length make it *look as though* the lines are the same length? (1984: 34).

If Shmax's prior outlook includes his knowledge of the MLI's true nature — or, in Fodor's words if "perceptual problem solving has access to ALL (or, anyhow, arbitrarily much) of the background information at the perceiver's disposal" (1984: 35) — then perceptual conservatism gives no clear explanation as to why Shmax's *experience* is not altered in addition to his *judgment*. That is to say, Shmax's knowledge of the MLI seemingly *could* hijack his experience of the MLI (in the sense that his knowledge could alter his experience), *but it doesn't*. If Shmax has robustly developed his understanding of the illusion such that <MLI> is now a part of his outlook, then perceptual conservatism (which contends that outlooks can alter experiences) is at pains to explain why Shmax's prior outlook does *not* make the lines of the MLI *look* to be the same length to Shmax.

Recall that the only comment Siegel (2017) gives to the MLI is that "If you know the lines are not the way they appear, and you cease to rely on the experience, then there is no further rational adjustment to be made in the situation" (36). There seem to be three ways to interpret what Siegel means by "cease to rely on" here:

- i) She is conflating <fixation of perceptual belief> with <perceptual experience>.
- ii) She is including <fixation of perceptual belief> as a step within the purview of <perceptual experience>.
- iii) She is referring to perceptual judgment.

The discussion in this section clearly discounts option (i). Option (ii) would require positing an additional step between a subject having an experience and a subject rendering a judgment about

that experience in a manner that would seem unavoidably guilty of special pleading. But option (iii) fails to address the recalcitrance of the illusory *experience* (not judgment) of the lines.

In contrast to Siegel, a perceptual dogmatist could argue that the phenomenology of (1) is sufficient (in the absence of defeaters) to justify (2): Max is *not* irrational (because he lacks defeaters) and neither is Shmax irrational (because he rejects (2) on the basis of his other beliefs).⁴¹ In a different way, a reliabilist could say that the *etiology* of (1) is sufficient (in the absence of defeaters) for (1) to serve as an epistemic foundation for (2): in the same way and for the same reasons as in the dogmatist's explanation, neither Max nor Shmax are irrational. But Siegel's attempt to rely on both phenomenology and etiology within the mechanisms of perceptual conservatism (and apply them directly to perceptual *experience* instead of judgment/perceptual-belief-fixation) makes for an unwieldy machine: recall that perceptual conservatism as described contends that an experience's "phenomenal character suffices to give experience baseline positive charge, absent defeaters or epistemically bad-making etiologies" (Siegel 2017: 47). On this model, Max's experience of (1) has a phenomenal character, was formed in a normal manner (without the influence of his outlook), and, because he lacks defeaters for it, has positive epistemic charge, so when Max uses (1) to form a belief in (2), Max is rational. Shmax, on the other hand, experiences the same phenomenal character described by (1) via the same normal manner (without the influence of his outlook), but his knowledge of the MLI is a defeater for (1), so (1) does *not* have positive epistemic charge for Shmax and he must "cease to rely" on (1) in order to be rational by rejecting (2). This is, I take it, the thrust of Siegel's explanation of "cease to rely" as captured by (iii).

⁴¹ Tucker (2013: 2) defines perceptual dogmatism as "if it *perceptually* seems to S that P, then, in the absence of defeaters, S thereby has justification to believe P."

But it is not clear how the defeater Shmax possesses (or *defeaters*, since any of the multiple truths Shmax knows about the MLI could serve as a relevant defeater in this case) could function in this way. For both dogmatists and reliabilists, the defeater serves to allow Shmax to consciously reject the move from (1) to (2) — if not for the defeater, Shmax would be just like Max and unwittingly make that move without thinking twice; for the perceptual conservative, Shmax possessing the defeater directly changes the epistemic charge of (1) without changing the phenomenal character, given that Shmax *still has to consciously reject* (2) — this is exceedingly strange if it is the phenomenal character which is supposed to give (1) its epistemic charge in the first place. Siegel might argue that Shmax doesn't need to be conscious of the defeater if it's just an unconscious part of his outlook,⁴² but this would still mean that Shmax's perceptual systems (which can, by hypothesis, interact with Shmax's outlook) are ignoring relevant information about the MLI that Shmax possesses in rendering an experience with erroneous presentational character. Either way, if perceptual systems are vulnerable to outlooks in the way that perceptual conservatism claims, then the defender of such a model carries the burden of proof to explain why outlooks *don't* penetrate perceptual experiences in the case of known illusions, but *do* penetrate experiences in other cases (such as the racist ones mentioned previously).

In sum, if perceptual experience can be rational as Siegel claims, then Shmax's experience of (1) is either rational or irrational. If (1) is irrational (in that it does not have the ability to empower beliefs), then Shmax should cease to rely on it and reject (2), just as Siegel asserts. But, if Shmax's outlook can also hijack his experience, then it remains a mystery why his outlook has not altered his experience of (1).

⁴² Siegel is explicit that “experiences are epistemically charged because they belong to one’s outlook on the world” (2017: 49).

Because much of Siegel's project is concerned with analyzing hijacked experiences that ultimately result in false beliefs, she has little to say about hijacked experiences that could (if they obtained) result in true beliefs. Similarly, the discussion of perceptual conservatism in *The Rationality of Perception* gives no consideration to the possibility that perceptual hijacking *could* be epistemically beneficial — it simply defines hijacking as pernicious (in much the same way that cognitive penetration is often treated).⁴³ Furthermore, it is not clear that, on its own lights, perceptual conservatism *can* delineate between epistemically beneficial and epistemically pernicious forms of perceptual hijacking; it is to this problem I now turn.

3. On Perceptual Hijacking

The case of Shmax and the potential boon perceptual hijacking could offer to perceptual experience mirrors Lyons' point that cognitive penetrability is not necessarily epistemically problematic *if* said penetration promotes the fixation of otherwise-justified beliefs; consider the following case from Lyons (2016a):

SNAKES: You and I are out hiking in the woods, and I believe (for no good reason) that these woods are full of snakes. This belief primes my visual system for snakes, enabling me to see through their camouflage, and it does so in a way that increases my accuracy without producing any more false positives. I end up visually spotting snakes that you don't see but that are really there in the woods around us (7).

Even though, by hypothesis, the unjustified initial belief about snakes penetrates the subject's perception, if this makes a *pro tanto* contribution to the subject's ability to detect real snakes in his vicinity, then those perceptual beliefs about nearby snakes would (as a function of their formation by a generally-reliable perceptual process) remain justified. In sum, Lyons explains

⁴³ For more on this analogy, see the next section.

that “the good kinds of cognitive penetration are the kinds that increase reliability, while the bad ones are the ones that decrease it” (2011: 300).⁴⁴

In a similar way, Shmax’s beliefs about the true nature of the MLI should be embedded within his prior outlook, so *if* those beliefs were to affect his perceptual experiences in the way that Siegel argues is possible, then the presentational phenomenal character of Shmax’s next encounter with the MLI should change, thereby altering the epistemic charge of Shmax’s perceptual experience. Given the role of Shmax’s prior outlook on this alteration, this would qualify as a case of what Siegel calls perceptual hijacking, but it would *not* be a pernicious one (insofar as Shmax’s beliefs would render him immune to the illusion). And, notably, this is true regardless of whether perceptual conservatism cashes out the role of Shmax’s outlook in terms of cognitive penetration or perceptual learning; because it is the *outlook* that carries the explanatory weight, whether the outlook’s penetration of perception is a diachronic or synchronic process is beside the point.

However, the MLI is famously *not* susceptible to such experiential alterations.⁴⁵ Perceptual systems theory can neatly account for this recalcitrance by explaining how various features of the lines simply trigger the outputs of certain percepts, regardless of any doxastic states *or* any subjective phenomenology. Because PST treats perceptual outputs as nonrational foundations for epistemic processing (or, on Lyons’ view, as basic beliefs that lack rich phenomenology), their presentational character is epistemically irrelevant. Both Max and Shmax see the same thing, but Shmax’s beliefs about the MLI preclude him from inferring that the lines

⁴⁴ Siegel (2017: 124) discusses SNAKES as a potential criticism of the inferentialism on which her RP thesis is based, arguing that the case is better explained as a modulation of the perceiver’s *attention*, rather than specifically their perceptual experience. Even if that analysis were accurate, the MLI case offers no such scapegoat.

⁴⁵ Although I remain frustratingly unclear about the semantics of the phrase “the exception that proves the rule,” Brandom (2015: 119) seems to function as such here.

are of different lengths *even though* the outputs of his perceptual systems nevertheless make it *look* to him like the lines are of different lengths.

Perceptual conservatism's grounding of an experience's epistemic charge in the phenomenology of such outputs precludes such an explanation: Max and Shmax have the same experience of the MLI, but it would seem that Shmax interprets or understands this experience in a different way than Max. If this is so, then something *other* than the simple presentational character of the experience is what grounds Shmax's beliefs, but perceptual conservatism gives no clear explanation for what this additional factor could be such that known illusions remain recalcitrant in the final analysis.

For example, Siegel suggests that *expertise* might qualify as the differentiating factor between Max and Shmax insofar as Shmax's skilled outlook allows him to understand the presentational character of (1) in a manner that Max cannot; in fact, Siegel argues that expertise could inferentially-modulate experiences to *upgrade* their epistemic power (2017: 144). However, in the particular case of the MLI, *no* amount of expertise seems sufficient to manifest a change in perceptual experience. Siegel (2017: 36) can argue that the MLI-expert could "cease to rely on the experience" of the MLI at which point there is "no further rational adjustment to be made in the situation," but that is beside the point: Shmax shouldn't have to *cease* to rely on his experience because *his experience should change*. Barring this, perceptual conservatism should, at the very least, provide some explanation as to why the interplay between his outlook and his perceptual systems has failed to guide the skilled-MLI-perceiver Shmax towards the truth as quickly as it (by hypothesis) could.⁴⁶

⁴⁶ Ghijsen (2018: 504) suggests that it is the "assertoric force of perceptual experience, the fact that it represents the world as being a certain way" is what truly matters here, but this is something that both Max and Shmax share.

If prior outlooks can affect perceptual experience (or “appearances,” to use Fodor’s language) in addition to perceptual judgment (“fixation of perceptual belief”), then this effectively jettisons the foundational starting point needed to prescind the justificatory regress, once more underlining the insufficiency of perceptual experiences *qua* perceptual experiences for grounding justified beliefs. Siegel seems to think that phenomenal grounds can play the foundational role on their own — insofar as such grounds can “provide their own epistemic charge” (2017: 51) — but this is simply ad hoc. No clear reason has been given as to why phenomenology could play the foundational role that Siegel’s argument needs it to play.⁴⁷

4. Why Perception is Nonrational

Recall, then, the Sellarsian dilemma, summarized by Lyons (2016b: 1062) as so:

1. If experiences lack conceptual/propositional/cognitive content, then they can’t justify beliefs.
2. If they have conceptual/propositional/cognitive content, then they can justify beliefs, but only if they are themselves already justified.
3. Either way, therefore, experiences cannot both terminate the justificatory regress and justify beliefs at the same time.

In Section II,⁴⁸ I argued that Siegel’s perceptual conservatism essentially grasps the experientialist horn of the Sellarsian dilemma — Premise 2 here — to assert that perceptual experiences can, under the right conditions, justify beliefs. Siegel’s model of epistemic charge purports to explain what it would mean for experiences (and/or their contents) to be “already justified” and treats phenomenal states as potential foundations for epistemic empowerment.

However, in addition to the seemingly intractable problems for phenomenology’s epistemic capabilities suggested by the Müller-Lyer illusion, there is no clear reason to think that

⁴⁷ See also Ghijssen (2018: 504–505).

⁴⁸ Specifically, Section II.2.

presentational phenomenal character can possess self-justifying epistemic power — that is, there is no reason to think that phenomenology, however qualified or constrained, *can* serve as an epistemic foundation. Epistemic charge, then, seems to be something like the fabled “luminiferous aether” once posited as the medium through which electromagnetic energy might travel;⁴⁹ it is an ad hoc definition of a concept needed to fill-in an otherwise problematic theoretical lacuna. As Lyons (2016b) puts it, any attempt to grasp the experientialist horn simply cannot be “...at all plausible unless we can somehow understand something literally serving as...evidence for itself” (1063). But even if this isn’t the case, it’s not clear why Siegel needs to accept Premise 1 (of Lyons’ construction), as she seems to do.

Lyons (2016b) points out that it seems easily conceivable for a creature to have the “brute awareness” of a headache while lacking the concept ‘headache’ (1063). If this is true, then the creature possesses a seeming of a headache (a) that is distinct from both the sensation of the headache (b) and the belief that “I have a headache” (c): the creature under consideration possesses both (a) and (b), but lacks (c) and the seeming qualifies as conceptual content but not as a belief, so (1) — and the complicated apparatus needed to undergird Siegel’s RP-thesis — can simply be rejected (1073–1074). Instead, Lyons proposes an essentially-Fodorian response: if perceptual beliefs are the output of a perceptual module, then they bypass the concerns of the Sellarsian dilemma altogether, given their grounding in a clearly non-epistemic fact: perceptual systems theory (PST). Moreover, they bypass the concerns that motivate Siegel as well: instead of aiming to defend the rationality of perceptual experience, perceptual system theory (and its

⁴⁹ For a brief overview of how the concept of aether was deployed in nineteenth-century physics, see Chalmers (2008: 163–164).

concomitant support for so-called “zombie epistemology”) can treat all such perceptual experience as a coincidental byproduct of modular operation (or something else entirely).⁵⁰

Both perceptual systems theory and perceptual conservatism offer explanations for the process by which perceptual belief–fixation operates, but only the former avoids or gives satisfactory answers to the questions raised in Sections II and III of this project. Insofar as perceptual conservatism remains committed to Siegel’s RP-thesis — “Both perceptual experiences and the processes by which they arise can be rational or irrational” (2017: 15) — it leaves itself open to the criticisms and concerns here described. By consideration of heuristic parity alone, PST is to be preferred, the RP thesis is to be denied, and perception is to be seen as squarely nonrational.

So, recall the cases of PINE and SCOPE, with which this project began. In PINE, Calvin gradually gains a recognitional disposition to distinguish pine trees from non-pine trees. In SCOPE, Thomas has a perceptual experience of a tiny embryo inside of a sperm cell (despite the fact that no such embryo exists).

In both of these cases, perceptual conservatism could argue that Calvin and Thomas’ distinct prior outlooks infect their perceptual processes, thereby altering their perceptual experiences in ways that either lead them towards or away from the truth. However, as argued in this Section, perceptual conservatism simply assumes that the phenomenal character of their experiences can play this epistemic role, thereby confusing explanations of why Calvin’s experience is upgraded while Thomas’ is downgraded. By treating perception as a directly

⁵⁰ Notably, although PST is a far cry from the model of perception that Siegel ends up defending, several of the defenses she poses of her RP-thesis seem reminiscent of the squarely reliabilist PST. For example, her treatment of inference-without-reckoning (2017: 77–106) seems highly similar to how modularity-defenders describe the output of perceptual modules as “inferentially opaque” (Lyons 2009: 95). Whereas Siegel insists that such an inference carries no particular qualitative character — it “doesn’t have to feel like anything” (2017: 82) — Lyons will simply call it “cognitively spontaneous in Bonjour’s sense” (2009: 95); either way, though the terminology is different, the process seems suspiciously similar in its automaticity.

rational process, instead of as a nonrational process which provides inputs for rational processing, perceptual conservatism muddies otherwise clear epistemic waters.

In contrast, Perceptual Systems Theory can explain PINE without reference to Calvin's perceptual experience at all: if perceptual *learning* is responsible for Calvin's improved process of belief-fixation regarding pine trees, then as Calvin's perceptual systems are gradually trained to output percepts, thereby triggering/grounding perceptual beliefs about pine trees, Calvin's recognitional disposition will likewise develop. Once again, Calvin's visual phenomenology — the seemings to which he has conscious access — is an experiential bonus to this otherwise-unconscious process of percept-generation and belief-fixation.

However, Perceptual Systems Theory cannot easily explain SCOPE. The most natural response would be to simply deny the case: Thomas *thinks* he sees an embryo on his microscope slide — he forms the perceptual *belief* that “there is an embryo” — but, by hypothesis, this leaves the perceptual experience of an embryo unexplained. Without an embryo present to trigger Thomas' visual perceptual system to output embryo-percepts, it is not clear why he would have a *seeming* (or a “perceptual experience”) of an embryo.

In Section IV, I consider this problem and the role of Thomas' *perspective* on the operation of both his perceptual systems and on the epistemic uptake of their produced percepts.

IV On Perspectival Epistemology

Despite perceptual conservatism's problems, its motivations are straightforwardly clear: the phenomenology of perceptual experience might, in at least some noteworthy cases, seem to be affected by higher-order mental states.

Consider the evidence surrounding the phenomenon of *weapon bias*, whereby a subject reports to have had a perceptual experience of a gun when, in reality, they were shown an image

of a hand tool; consistently, across a wide range of studies, test subjects primed with images of black faces are far more likely to experience weapon bias than those primed with images of white faces (Kubota and Ito 2014; Jones and Fazio 2010; Fleming, Bandy, and Kimble 2008; Payne 2006). A standard explanation for such findings is that implicit racist attitudes unconsciously penetrate and thereby precipitate biased experiences and reactions (Mandelbaum 2016; Gendler 2011).⁵¹

Consider the case of Stephon Clark, who was killed by two members of the Sacramento Police Department in his grandmother's backyard; after firing twenty rounds and hitting Clark seven times, both officers swore that they believed Clark — who was Black — to be holding a gun, despite the fact that only a cell phone was found on Clark's dead body (Wamsley 2018).⁵² Siegel's perceptual conservatism gives a simple explanation for this case: the officers' prior outlooks (which could include all manner of racist beliefs, preferences, affective responses, and the like) infected their perceptual experiences such that the content of those experiences differed, but said experiences — though phenomenologically salient to the officers — were epistemically *downgraded*, so the officers' corresponding perceptual beliefs (something like “He has a gun”) were unjustified.

Siegel is right to note that these sorts of perceptual experiences are as problematic as their reports are commonplace, but her chosen explanatory route — treating perceptual experience as a basic epistemic unit which can be rationally assessed — raises more questions than it answers. Instead, I propose that defenders of perception's modularity can adapt the model of cognitive architecture defended by Elisabeth Camp to solve the same sorts of perspectival problems of

⁵¹ The exact nature of these implicit attitudes, however, remains a lively debate, see Brownstein and Saul (2016a; 2016b) for an introduction.

⁵² Although the officers responsible for Clark's death were not formally charged with a crime (Adler 2019), the city of Sacramento agreed to pay a settlement of \$2.4 million to his children (Vigdor 2019).

perception without sacrificing the virtues of a modular view like perceptual systems theory. In short, racist police officers — as well as preformationists like Thomas in SCOPE — have developed perspectival dispositions which trigger their visual perceptual systems to output percepts formatted as what Camp calls *characterizations*. These perceptual outputs are not directly rational, but undergo epistemic uptake via the already-discussed process of perceptual belief-fixation. So, the officers who shot Stephon Clark might well have had a perceptual experience as-of a gun in Clark’s hand, insofar as they had dispositions to perceptually characterize Black men as threats, thereby having shaped their perceptual systems to output racistly-characterized percepts — but the percepts and the experience are, as discussed in Section II, logically distinct.

To flesh out this explanation in more detail, I first briefly outline Camp’s model, highlighting the portions relevant for the discussion presented here. Then I consider evidence for categorical perception and the so-called “shallow” outputs of modular perceptual systems. Finally, I demonstrate how, if the so-called “Camp Doctrine”⁵³ is correct, then the perception of categories could indeed subsume characterization-talk in a manner conducive to explaining cases like Stephon Clark and SCOPE without needing to reject phenomenological explanations wholesale. Furthermore, because it makes no particular commitments about epistemic justification (or anything like ‘epistemic charge’), I argue that a Camp-ish explanation is preferable to the tenuous model of perceptual conservatism.

⁵³ This tongue-in-cheek description of Camp’s tripartite model of characterizations, perspectives, and frames originates from several public talks and presentations she has given, as well as in lighthearted personal conversation with the author.

1. Characterizations and Concepts

Before I can properly characterize the ontology of a perceptual system's outputs, Elisabeth Camp's distinction between logical concepts and associative characterizations must be briefly outlined. Although philosophers are apt to treat concepts as the fundamental atoms of propositional thinking, Camp argues that they are only one kind of cognitive representation serviceable as inputs for higher-order thought: equally basic are associative clusters she calls characterizations.

a. What Characterizations Are

Drawing from well-established positions in both cognitive science and representationalist theories of philosophy of mind (Carey 2009; Fodor 1983), Camp treats concepts as “word-like representational structures that function as stable, arbitrary, recombinable bits,” able to be deployed in a multiplicity of contexts without changing (2015a: 591). Accordingly, a concept's representational format is as a discrete atomic entity that, by serving the same referential role systematically, satisfies what Evans (1982: 104) dubbed the “generality constraint,” insofar as it is recombinable within a seemingly unlimited number of propositions (Camp 2004: 210).⁵⁴ Ultimately, concepts serve to “bring together multiple instances as belonging to the same kind, either by ascribing a common property to multiple objects or by re-identifying a single object as it gains and loses properties” (2015a: 593).

In contrast, characterizations are “intuitive, holistic, and context sensitive” representational structures that dispose a person to “*gestalt* subjects in certain ways” (Camp 2015a: 601). Similar to how cognitive psychologists discuss concepts as an “open system” of

⁵⁴ One description of the constraint runs: “We thus see the thought that *a* is F as lying at the intersection of two series of thoughts: on the one hand, the series of thoughts that *a* is F, that *b* is F, that *c* is F,..., and, on the other hand, the series of thoughts that *a* is F, that *a* is G, that *a* is H” (Evans 1982: 104, n. 21).

contextually-dependent bridges between minds and the world (Rosch 1999: 61), characterizations are “informationally, experientially, and affectively rich, integrating as much data as possible into an intuitive whole” (2019: 20).⁵⁵ Similar to Rosch’s theory of prototypes, which explains mental categorization by positing super- and sub-ordinate levels of mental representational files organized in contextually defined arrangements as specified by a subject’s history (1978), characterizations are complex webs of multiple properties that are inextricably knit together and accessed only by context-sensitive triggers. Characterizations are expressly *not* propositional in their format, consisting instead of a hybrid-like structure of both iconic and discursive components arranged in particular patterns or hierarchies, similar to a cartographic design (Camp 2018a, 2015b).⁵⁶ Consequently, a token characterization is *not* freely recombinable within propositions, but rather implicates a particular batch of properties in association with a particular stimulus whenever that characterization is mentally represented.⁵⁷

Camp gives the example of a football team’s quarterback:

my characterization of quarterbacks includes their being natural leaders, affable, and a bit shallow. In addition to such general traits, characterizations also often include more specific, experientially represented properties: thus, I think of quarterbacks as having a certain sort of square, clean-shaven jaw, gleaming teeth, and a ready smile (2015a: 603).

Some of the attributional features packaged into a characterization are similar to concepts, but many of them (like a “way of carrying oneself,” the qualitative/experiential character of vocal

⁵⁵ Indeed, Camp explicitly suggests that typical disagreements between philosophers and psychologists about the nature of concepts might be cashed out as a semantic distinction between concepts (in the philosophical sense) and characterizations (in the psychological sense of “concept”); see Camp 2019 (19, n. 1). For a similar analysis, see Carey (2010: 490).

⁵⁶ They also bear a marked similarity to the complex role of “content” Johnston (2006) discusses (complete with scare quotes): “The “content” of sensory awareness is not a propositional content, but is rather a host of interconnected exemplifications of properties, relations, and kinds. That “content” is *pre-predicative* in the sense that awareness of it need involve no predication, no judgement that such and such is the case, nor any proto-judgment implicitly deploying the concepts on which the corresponding explicit judgment depends” (282–283).

⁵⁷ Notably, Camp indicates that characterizations can apply to both types and tokens: just as I can have a characterization about “American presidents,” so too can I have a characterization of “Barack Obama” (604).

tone, or an affective mood provoked by a subject's presence) are only referenceable via demonstratives or semantic tools like metaphor (Camp 2015a: 603): Peacocke's protopositional contents (1992: 77), Burge's perceptual attributives (2010:36), Sellarsian "this-suches" (1968: 5), and the primitive semantic values conversely described as both propositional radicals and propositional skeletons (Cappelen 2007: 207) might all be candidates for the sorts of representational features bundled together in one of Camp's characterizations. Importantly, characterizations also incorporate affective valences which, because characterizations are *not* reducible to any single component of the package, serve to further underline individual characterizations' relegation to subjective, often idiosyncratic judgments.

However, to call a characterization a 'judgment' would elide two further important distinctions between these associative webs and digital concepts: the lack of an endorsement constraint and their holistic structure. The former indicates that a subject who mentally represents a token characterization about an extramental object is not required to believe that the object *actually* possesses the features attributed to them by the characterization; it is simply the case that, insofar as the represented object is a member of the class picked out by the characterization, then the perceiving subject "take[s] those features to be *fitting* for them" (604). Of course, *fittingness* is orthogonal to *truth*, a fact Camp explains by observing that "Although it might be nice if fittingness could be straightforwardly reduced to statistical norms, intuitions of fittingness often appear to have a more squarely aesthetic basis...If we were more fully rational, we would sharply distinguish what we take to be fitting from what we believe to be actual or even probable" (604–605). Similar to the distinction Siegel makes between a complete and a considered outlook, Camp's characterizations are automatically generated by a subject, who can

then consciously affirm or deny subsets of the package, but the full characterization is mentally represented nonetheless.

The final feature of characterizations which sets them apart from concepts is their holistic structure, an important element also for underwriting their viability as an architectural element of cognition. Whereas concepts have a minimal internal structure, abstracting all extraneous details away from the discrete conceptual essence, characterizations arrange the panoply of represented features in a multidimensional pattern analyzable in terms of both *prominence* and *centrality* (2015a: 605). Prominence is a measurement of how apt a particular feature is to activate the representation of the characterization as such,⁵⁸ while centrality is an appraisal of how significant a given feature is within the characterization as a whole; as Camp explains, “Where prominence selects which features matter, centrality determines how they matter, by connecting features into explanatory networks, such that more central features are more richly connected to other features” (2019: 20). Note that, ideally, this is more significant than distinctions between so-called “thick” and “thin” concepts (Väyrynen 2016); the complexity of even the most threadbare characterizations, particularly in terms of the hierarchical design of the internal structure of their features, are meant to outstrip the thickest sort of proper concepts.

Altogether, the aesthetic nature of characterizations, and their assessability *vis-à-vis* fittingness and dimensions of prominence and centrality, not only draws natural comparisons to aesthetic judgments in general, but also suggests an evolutionary advantage to mental representations with the vibrancy and complexity of Camp’s characterizations. Consider the perceptual gestalt in [Fig. 4] which, depending on how a subject views it, might appear to be

⁵⁸ Prominence is further analyzable in terms of *diagnosticity* (a measure of how useful a given feature is for provoking object categorization) and *intensity* (a measure of a feature’s signal-to-noise ratio compared to other features in the package).

either an elderly woman facing the bottom-left corner of the image or a young woman with her head turned away from the viewer.



*Fig. 4: “My Wife and My Mother-in-Law”*⁵⁹

Camp points out that while it is one sort of experience to sequentially process which portion of the drawing is supposed to be the old woman’s nose or the young woman’s ear, it is an entirely different sort of experience to actually see the nose *as* a nose or the ear *as* an ear. Moreover, this cannot be done by halves: while a viewer does have attentional control over which way the image is perceived, a viewer cannot see the image halfway as an old woman and halfway as a young woman simultaneously: “ultimately the “click” of holistic understanding is something that just happens — or doesn’t” (Camp 2015a: 610). When faced with a situation that could be dangerous or extremely time-sensitive, the immediacy of information provided to a subject by her characterizations could be crucial, even if that information is ultimately less than perfectly accurate (i.e. includes unendorsed features, for example); characterizations need not always track

⁵⁹ First published by W. E. Hill in 1915, now available at <https://www.loc.gov/pictures/item/2010652001/>.

the truth, provided that they do so more frequently than they don't (that is: provided that they, in general, 'fit' their relevant contexts).

b. What Characterizations Are Not

On Camp's model, whereas concepts are something like atomistic mental representations with discrete, context-neutral semantic content, characterizations are basic mental representations with adamantine clusters of features that provide context-dependent semantic content. Though this bears similarities to dual-factor theories of concepts that divide conceptual content into *wide* (or *broad*) and *narrow* aspects, it is importantly different. Briefly, the former is determined by ostentional factors fixed by a subject's environment and the latter is not,⁶⁰ so a characterization's context-sensitivity might make it seem similar to wide factors of concepts, but this assessment is incomplete. In Camp's model, characterizations might incorporate features that are most conveniently indicated demonstratively, but this is not to say that their content is exhausted by such features; indeed, insofar as much of a token characterization's internal structure incorporates first-person subjective elements (like emotive responses), it incorporates paradigmatic narrow content as well. So, the concept/characterization distinction is simply orthogonal to that of wide and narrow content; if arguments meant to distinguish between wide and narrow conceptual content succeed, then it might well be the case that Camp's characterizations are similarly divisible as well.⁶¹ Put differently, if the meanings of concepts "just ain't in the head," then similar arguments might be run for *some* (but not all) elements within characterizations, but it would be incorrect to identify Camp's concepts as "what's in the head" and characterizations as "what's not" or vice-versa.⁶²

⁶⁰ On this see Carey (2011: 161–162) and (2010: 514–523); Mandelbaum (2011: 138); and Block (1987)

⁶¹ Developing this point is beyond the scope of this thesis.

⁶² See also Putnam (1975: 144) and Devitt (1990).

In roughly the same way and for roughly the same reasons, characterizations are not simply Gibsonian affordances (1979): they are internal mental representations, not external-world properties. One recent development of affordance-based philosophy of mind is Ruth Millikan’s architectural model of *unitrackers* and *unicepts* designed to supplant talk of concepts entirely (2017). For Millikan, a unitracker is a cognitive mechanism for tracking discrete things in our experience (in particular, though not exclusively, affordances) in a manner similar to how a visual index (or FINST) “binds visual objects to their representations” (Pylyshyn 2003: 202).⁶³ Conversely, a unicept is the collection of “stored information about the same thing together” which a unitracker targets when it performs its binding function (Millikan 2017: 7);⁶⁴ in contrasting them with classical concepts, Millikan suggests that unicepts “may not be sharply separated units, but might be overlapping or continuous networks” (48). Although it might seem like the discrete focus of unitrackers and the complex network of unicepts could be mapped onto the atomic structure of Camp’s concepts and the associative web of characterizations, this is a panglossian oversimplification. For example, Millikan is adamant that both unitrackers and unicepts are *particular* representations that different individuals do not share,⁶⁵ whereas Camp treats concepts as standard universals that agents represent and characterizations as representations that, though often idiosyncratic to agents, can be shared insofar as two agents share a perspective. In short, Millikan’s model concerns competencies or capacities that agents

⁶³ Unitrackers are more robust than FINSTs, though, given that they serve to track abstract and immaterial objects (like <intelligence> or <democracy>) as well as objects detectable by visual perceptual systems.

⁶⁴ For more on visual indexes, see the material on object files referenced in Note 79.

⁶⁵ As she explains, “[i]t is not unitrackers or unicepts that are “shared” but only their referents or extensions — things outside the mind and not within” (47). So, whereas it makes sense to say “X and Y have the same <electron>,” it would not make sense to say “X and Y have the same electron-unicept,” but rather something like “X’s electron-unicept is functionally identical to Y’s electron-unicept.”

possess and develop for navigating *the* world;⁶⁶ Camp’s model concerns the building blocks by which an agent comes to perceive *their* world.

Millikan’s notion of unicepts intentionally develops her long-standing notion of *pushmi-pullyu representations* (PPRs); intended initially to primarily explain behavioral responses to stimuli, PPRs are mental representations with both informational and imperatival content (Millikan 2004: 77–78).⁶⁷ If a fox in a forest hears a noise, it might mentally represent that noise as a PPR which provokes it to hide from potential danger; such representations could explain, then, how an animal (presumably) incapable of higher-order propositional thought could, nevertheless, respond to environmental stimuli and safely navigate their world. But even thinly developed characterizations are more complex than simple input-output mechanisms and they do not necessitate certain behavioral responses; my characterization of a quarterback does not entail a behavioral response — even if such reflexive components are potentially affixable within a characterization.

In this way, and for the same reasons, Camp’s characterizations are also distinguishable from aliefs. First posited by Tamar Gendler (2008a), an alief is an “innate or habitual propensity to respond to an apparent stimulus in a particular way” (553). Unlike beliefs, aliefs are *automatic* responses triggered by *associations* of mental events and environmental stimuli that function *arationally* and often include *affect-laden* and *action-generating* components; importantly, Gendler emphasizes that aliefs are primitive attitudes and can function as grounds for higher-order states like beliefs or desires (558). In Gendler’s original model, aliefic content is necessarily tripartite, as she explains:

A paradigmatic alief is a mental state with associatively linked content that is representational, affective and behavioral, and that is activated—consciously or

⁶⁶ For another capacity-based model of perception, see Schellenberg (2018).

⁶⁷ See also Millikan (2017: 64–65) for a specific comparison of PPRs and unicepts.

nonconsciously—by features of the subject’s internal or ambient environment (2008b: 642).

The representational-affective-behavioral (R-A-B) structure of aliefs provokes a cluster of specific motor and emotional responses to specific mental events: under this kind of model, if the agent is consciously aware of the representations and possesses intentional control over the responses, then they are beliefs; if he is not conscious and not in control, then they are aliefs (Kriegel 2010: 476). But, although characterizations do include mentally represented features that might escape an agent’s conscious awareness and, like aliefs, lack endorsement constraints, they are further unconstrained by the rigid aliefic R-A-B framework (or any necessary navigational or experientially sequential algorithms) and do not require affective or behavioral components.⁶⁸

Finally, it would be oversimplified to treat characterizations as reified representations of Fregean mental content (with concepts treated similarly as Russellian contents); given that characterizations are meant to incorporate semantic values alongside affective, behavioral, and other featured components, they are more variegated and complex than simple *de re* referents crystallized into a bundled representation.

So, although the distinction between concepts and characterizations here described bears a family resemblance to several recent models of associative representations, Camp offers a distinct explanatory framework for understanding the fundamental components of cognition.

c. The Camp Doctrine

Ultimately, Camp’s notion of characterizations serves to underwrite her perspectival model more generally. Although characterizations are idiosyncratic to individuals, agents have

⁶⁸ For an adapted model of implicit attitudes that offers a less rigid framework for their operation, see Brownstein (2018), Brownstein and Madva (2012a; 2012b). Importantly, even Brownstein and Madva’s FTBA model of aliefs is too formulaic for Camp’s concerns.

consistent tendencies to characterize given stimuli in consistent ways; this propensity Camp calls a *perspective*, or “an open-ended disposition to characterize: to encounter, interpret, and respond to some parts of the world in certain ways” (2019: 24). Insofar as agents are members of social groups (or share various identity-shaping traits with other individuals) then they may share perspectives (to at least some degree); consider members of the same religious group who share a significant number of overlapping dispositions to characterize relevant things in a manner typical of an “Evangelical Protestant” (Camp 2019: 25). Finally, communicating one’s dispositional tendencies, or provoking another to experience the aforementioned “click” of holistic understanding is best accomplished via the deployment of a representational vehicle (such as a slogan or catchphrase) which Camp calls a *frame* (2020, 2019). Although characterizations have a robust structure and perspectives are commonly activated dispositions, because neither one has a *propositional* structure, they are not easily communicable in full; the propositional format of frames (or their ability to be described propositionally) allows them to “crystallize perspectives in compact, explicit form” such that they can be communicated (2019: 27). Altogether, the tripartite “Camp Doctrine” (of *frames* which communicate *perspectives* which are dispositions to *characterize* objects in defined ways) offers a fruitful model of perspectival epistemology with many promising research projects in store.

The important thing for Camp’s model at this point is that *neither* concepts *nor* characterizations are more basic elements of our cognitive architecture; both serve valuable functions in different ways. Although concepts can be packaged into characterizations (alongside emotive or experiential elements), this relationship can also work in reverse, such as through the conceptual features inherent in the deployment of frames. Ultimately, though, both simple

concepts and complex characterizations function as basic building blocks for higher-order thinking.

2. Perceptual Characterization and the Contents of Perception

Camp's varied discussions and applications of characterizations, perspectives, and frames have been, to date, typically pragmatically motivated to addressing other topics;⁶⁹ there remains, as yet, no clear articulation of the full set of necessary and sufficient conditions for what constitutes a characterization nor for how an individual negotiates the process of characterizing a particular in a given way. And though I do not intend to provide a comprehensive heuristic in the limited space remaining, I would like to suggest that Camp's characterizations provide valuable conceptual material for accommodating the available empirical data about the contents of perception in a manner which allows for an explanation of SCOPE not grounded in visual phenomenology (a la perceptual conservatism).

Frequently in this work, I have referred to the "outputs of a perceptual module" and dubbed such outputs "percepts," but I have yet to explain what such a percept might be. In Section II.1, I contrasted percepts with sensations: whereas the latter is phenomenologically rich and, thereby, responsible for giving perceptual experience its presentational character, percepts are *nonexperiential* outputs of perceptual systems that have conceptual content. So, whereas sensations grant perceptual experience with its qualitative nature, percepts are able to underwrite beliefs.

One question in need of clarification is the requisite informational *depth* of the percepts output by visual perceptual systems; that is to say, how much (or *what kind*) of content can a single percept offer for epistemic uptake? On this, Fodor's initial descriptions of perceptual

⁶⁹ Camp (2020), for example, focuses on scientific methodology, whereas Camp (2018b) discusses slurs.

module outputs are foundational: he argues that “the more constrained the information that the outputs of perceptual systems are assumed to encode—the shallower their outputs, the more plausible it is that the computations that effect the encoding are encapsulated” (1983: 87); since he argues on other grounds that perceptual systems are encapsulated, Fodor concludes that their outputs must likewise be shallow. The most basic potential outputs, then, might be structural features of objects like color, shape, orientation, or other primitive physical cues: on this view, percepts of such distal properties would serve as input for conceptual categorization processes. When I perceive a tiger standing in front of me, my perceptual systems output a panoply of radically shallow percepts (for properties like “orange,” “striped,” and “cat-shaped”) which then combine conceptually to underwrite my belief that “There is a tiger standing in front of me.” This process of recognition that “the object in front of me” is a member of the class “tiger” is the process of categorization and takes place extramodularly on the conceptual side of the perception/cognition border.

But, notably, Fodor does *not* argue that outputs (what I have referred to as ‘percepts’) are simply structural features of the external environment: instead, Fodor offers *basically categorized* or “‘basic’ perceptual objects” — à la Rosch et al. (1976) — as an example what sort of outputs perceptual systems might render, saying:

basic categorizations are typically the most abstract members of their inferential hierarchies that *could* be assigned by an informationally encapsulated visual-input analyzer; more abstract categorizations are not reliably predicted by *visual* properties of the distal stimulus. And basic categorizations are the ones that you would want the input systems to deliver assuming that you are interested in maximizing the information per unit of perceptual integration (as, presumably, you are). So, the suggestion is that the visual-input system delivers basic categorizations (97).⁷⁰

⁷⁰ Fodor also argues on grounds on parity of reasoning that, at the very least, natural kinds are a strong candidate (93).

A “basic categorization” in this sense is the non-arbitrary termination of a classification hierarchy: consider the list *tiger, feline, mammal, animal, physical object, thing* — in such a nested taxonomy, *tiger* is the most basic entry.⁷¹

So, by identifying Roschian categorizations as potential contents of perception, Fodor appears to offer pre-existent support for something like the Rich Content View proposed by Siegel (2011) and developed by many others to expand the scope of perceptual content beyond simple, low-level properties. Typically, views that defend rich content also expressly place the locus of categorization to the perceptual side of the perception/cognition border: if my visual perceptual systems can output rich contents, then perceiving a tiger standing in front of me might provoke the production of a rich percept like “tiger.”⁷² One way of considering the Rich Content View is as a thesis that not only basic categories (like *tiger*) but subordinate categories (like *Siberian tiger* or *Caspian tiger*) are represented directly in perception.

Mandelbaum (2018) argues convincingly (and for architectural reasons) that perceptual categorization is a proper part of perception and demonstrates how modular theories, in particular, are well-served by such a thesis. By comparing empirical data for the speed of perceptual categorization to standard processing speeds of neural circuits subserving the visual system, Mandelbaum demonstrates how perceptual systems are able to render accurate category identifications before any signal could traverse the reentrant loop for vision; the “absurd rapidity of perceptual categorization” functions at speeds as quick as 13 milliseconds, whereas it takes at least 50 milliseconds for cognitive processing to access and receive information from the

⁷¹ This list is based on the one Fodor provides on (94), adapted herein for contextual reasons.

⁷² This also demonstrates, in part, how — although they are often discussed in tandem — discussions of rich content are orthogonal to questions of cognitive penetration: it is clearly not necessarily synchronic mental states alone which could provoke the production of rich contents.

modular visual system (2018: 271–274).⁷³ According to Mandelbaum, though, while the outputs of perception are category representations, they are only of *basic-level* categories (like *tiger*), not the richer contents of subordinate categories (like *Caspian tiger*).⁷⁴ These findings comport well with the empirical evidence about the ultra-rapid speed of gist perception: consistently, studies indicate that subjects represent basic — and even *superordinate* — Roschian categories more quickly than subordinate ones.⁷⁵

From an evolutionary standpoint, the adaptive benefits of ultra-rapid perceptual categorization are clear: if I am a creature in an environment where tigers are a threat, then the more quickly I can categorize a distal stimulus as “tiger,” the greater survival advantage I develop. It is far less important to distinguish between breeds of tiger (as indicated by subordinate categories) than it is to identify tigers as such, particularly for triggering survival responses (like flight behaviors).⁷⁶ And, presuming that there are no overly-maladaptive consequences of triggering survival responses unnecessarily, survival of the fittest would presumably select for creatures who categorize superordinate categories more quickly in the same way and for the same reasons.

However, if the upshot for ultra-rapid categorization processes is as described — and if behavioral or other reflexive responses are indeed part of the explanans for the process — then it seems like something more complicated than just atomistic concepts is necessary for

⁷³ As in Note 73, this further subserves the distinction between questions of perceptual content and etiological questions about how such content is generated (à la concerns about cognitive penetration).

⁷⁴ See also Potter et al. (2014), Proctor and Brosnan (2013). It is worth noting that these studies were not testing the potential role of expertise on perception, so their applicability to questions of perceptual learning are tenuous.

⁷⁵ Notably, Mandelbaum offhandedly (and incorrectly) lumps superordinate and subordinate categories together in his analysis (2018: 277), but this mistake does not undermine his fundamental argument about the ultra-rapid speed of basic categorization. In fact, the studies he references did not actually test superordinate categorization.

⁷⁶ Moreover, depending on the ecological context, there may not be more than one species of tiger which I might feasibly encounter in my environment.

underwriting such explanations. In a footnote, Mandelbaum admits that some of the test stimuli used in the study he references:

were not merely basic-level conceptual activations. For example, some of the stimuli depicted relations between basic level concepts (e.g., children holding hands, a bear catching fish, boxes of vegetables). So it at least appears that the output of perception need not be the activation of a simple concept, but can be complex, which opens up the possibility that it might in fact have propositional structure (2018: 275n15).

The notion that perception's contents are fully *propositional* (as opposed to merely conceptual) is controversial, but if our architectural model of cognition treats logical concepts as solitary prime elements, then it is difficult to imagine a different explanation for what "complex outputs of perception" might be.⁷⁷ If, instead, concepts and characterizations together comprise the foundation of the house of reason, then an alternative explanation is available: percepts could include semantic content encoded as *characterizations*.⁷⁸

In summarizing Camp's description of characterizations, I have already compared them favorably to Rosch's theory of prototypes: both are anisomorphic, informationally rich property clusters developed in response to an individual's ecology and history and tokened in response to context-dependent cues.⁷⁹ If Fodor, Mandelbaum, and others argue that basic Roschian categories can be perceptual outputs, then this suggests that other representations with sufficiently similar structures — such as Camp's characterizations — could be perceptual outputs as well. Importantly, because characterizations are sedimented conjunctions of features gradually formed as a result of an agent's interaction with her environment, their content is not subject to synchronic penetration: to recognize characterizations as at least some percepts is not

⁷⁷ I discussed Sellars' critique of perception having propositional content in Section II.2; for a more contemporary treatment see Gupta (2006).

⁷⁸ For another defense of pluralism for perceptual outputs, see Quilty-Dunn (Forthcoming). Although not identical with perceptual psychology's notion of object files, Camp's characterizations might also be fruitfully compared to the multiple-slot view of object files found in Green and Quilty-Dunn (2017). See also Green (2018).

⁷⁹ See Section IV.1.a.

to sneak a mechanism for cognitive penetration into the explanatory framework but, rather, is simply the admission of a format borne from a complex process of diachronic penetration (in conjunction with other cognitive elements). That is to say, modelling the outputs of perception as characterizations provides a more expansive explanatory framework for the discussion of perceptual learning as a whole.

In this way, the phenomenon of weapon bias precipitated by racial cues can be explained as a side-effect of malformed characterizations: if the police officers who killed Stephon Clark had, over the course of time, bundled perceptual attributives, protopositional contents, affective responses and the like together into characterizations of “Black men” which underwrote their belief that Clark was a threat, then they may well have felt justified in their use of lethal force. Camp is clear that, in themselves, characterizations are not truth-apt (or epistemically appraisable at all); however, the dispositions to characterize extramental particulars in given ways (and, by extension, the linguistic frames which can trigger such dispositions) certainly *are* (Camp Forthcoming; 2019: 31). Racist perspectives can be analyzed and criticized for their infelicitous etiologies, their moral failings, and — of particular importance for questions of perception — their unreliability for producing world-accurate characterizations (Forthcoming: §5.1): so, no matter how much the officers might have *felt* justified for attacking Clark, the informational inaptness of their perspectives (in addition to their moral repugnance) is fundamentally indefensible and their actions cannot be excused simply because “it *seemed* that way to them at the time.”

On a perceptual systems theory, seemings can confer *prima facie* justification on beliefs because they have been formed in virtue of percepts produced via the reliable process of a perceptual system; one worry about eroding the border between cognition and perception is that

such penetration introduces interference with this otherwise trustworthy process, potentially cancelling the reliability of the perceptual system.⁸⁰ Siegel’s perceptual conservatism obscures the sufficiency of reliability for conferring justification by incorporating the presentational character of the percepts into its operation and, therefore, needs to posit an additional heuristic — epistemic charge — to account for the unnecessary variable of phenomenology. In so doing, however, Siegel can explain how the police officers’ seemings were faulty: they had been hijacked by the officers’ racist outlooks such that their content was altered. What the defender of a straightforward perceptual systems theory requires to match this explanation (without exposing themselves to the weaknesses of perceptual conservatism discussed in Section III) is a way to explain the *unreliability* of the officers’ seemings without recourse to an alteration in content: this is the explanatory power of accommodating malformed characterizations and their ability to *erroneously gestalt* perceptual contents.

Recall that characterizations are complex collections of features that include things like perceptual attributives, protopositional contents, and affective responses. Crucially, characterizations are *nonpropositional*, instead implementing a holistic structure that either does or does not “click” into place; Camp (Forthcoming: §3) argues that a proposition can be believed, whereas an agent must “get” a characterization — in much the same way that a person “gets” that the person they see in [Fig. 4] is either an old woman or a young woman.⁸¹ Recognizing characterizations as mental representations that require implementation, then, underlines an appreciation of cognitive gestalts wherein, as Camp explains, “just as with perceptual gestalts, then, implementation makes a palpable difference *which cannot be reduced*

⁸⁰ This is why Lyons (2016a; 2011) can defend the utility of cognitive penetration in constrained cases where said penetration nevertheless reliably encourages the formation of true perceptual beliefs.

⁸¹ Importantly, it is possible to “get” a characterization without believing or affirming its content; Camp offers as an example of this how you might understand a slur without believing that the slur is accurate or appropriate (Forthcoming: §3).

to a difference in represented content” (Forthcoming: §3, emphasis added).⁸² While holding my gaze at [Fig. 4] constant, the contents of my perception do not change, but I can nevertheless modulate my experience of the ambiguous figure by willfully characterizing those unchanging contents in different ways. So, if I could somehow *wrongfully* gestalt the contents of my perception (whether via synchronic or diachronic processes), then I would see those contents as something incorrect, even though the contents of my perception would be accurate — I would effectively be misinterpreting a nonambiguous figure as if it *were* ambiguous.

So, consider two police officers: one with malformed, racist characterizations of Black men and one without. If either officer were to view a Black man holding a cell phone, their perceptual systems would reliably output the same percepts of the cell phone, but those percepts must be implemented by the individual in order to generate the seeming: the racist officer implements those percepts such that he sees the cell phone *as* a gun, whereas the non-racist officer sees it simply as a cell phone. Camp describes this kind of upshot for the implementation of characterizations in perceiving gestalts as so:

This difference [between ways of seeing-as] is not just phenomenological, but functional: the agent can now navigate the domain in a flexible, fluent way they could not in the absence of the gestalt; and they are disposed to impose different labels and to treat different alterations to the domain as enhancing or undermining its coherence than they would not under a different construal (Forthcoming: §3).⁸³

So, both officers act on the basis of their seemings, but the racist officer has tampered with the reliability of his perceptual systems, thereby undermining the justification that percepts output by those systems could confer. This would mean, then, that perceptual beliefs formed by the racist

⁸² Camp also quotes Wittgenstein (1953) here: “the agent ‘sees that it has not changed and yet sees it differently.’”

⁸³ Although Camp discusses some perspectival results of perceptual processes, she undertakes no explicit discussion of the nature of perception as such; the closest she comes is in a treatment of perspectival dispositions to *attend* to a visual scene in a certain way (Forthcoming: §2.1).

officer about what the man is holding will not only be false, but (given their formation in virtue of a malformed characterization) unjustified.

So, the case of Stephon Clark can highlight a key difference between Siegel's perceptual conservatism and a perceptual systems theory informed by Camp's characterizations: in the first case, the racist prior outlook of the police officer hijacks his perceptual experience, altering the contents of his perception such that he *genuinely experiences* a gun in Stephon Clark's hand (despite no such gun existing) and then forms unjustified beliefs on the basis of that ill-formed experience; in the second case, the racist perspective of the police officer predisposes his perceptual systems to output percepts that (inaccurately) characterize Clark as a threat — the officer might *think* (and even *report*) that the contents of his perceptual experience included a gun, but *they did not* — the officer was simply *wrongfully gestalting* the actual contents of his perceptual experience (a cell phone) while forming beliefs *vis a vis* his racistly-characterized percepts.

Finally, in the same way, this treatment of characterizations offers an explanation for SCOPE that does not require simply rejecting the subject's reports of his perceptual experience.

Recall the case:

SCOPE: Thomas is a well-informed biologist in the early seventeenth century who believes in preformationism (the theory that organisms develop from smaller versions of themselves). When Thomas inspects a sample with his microscope, he has a perceptual experience of a tiny embryo inside of a sperm cell.

Just as in the case of the racist perspectives of the police officers in the Clark case, Thomas has, over the course of time, developed a disposition to characterize things as a preformationist — that is to say, he has developed a preformationist perspective. When he sees the sperm cell on his slide, his perceptual systems render characterizations of the sperm cell in virtue of which he

forms the belief “There is an embryo on my slide,” but this does not entail that the contents of his perceptual experience are actually different as a result of his perspective. Instead, his aberrant perspective of the world leads him to inaccurately characterize the slide’s contents, both in his conscious perceptual experience and the process of perceptual belief-fixation. Having sedimented his perspective as a preformationist, he has, in effect, corrupted his perceptual systems — learned to see badly, as it were — so, insofar as his perceptual systems output malformed characterizations, the perceptual beliefs he forms in virtue of those characterizations (such as “There is an embryo on my slide”) are unjustified.

In short: Thomas’ perspective leads him to incorrectly gestalt, and thereby misunderstand, what he sees, but — because it results from how his perceptual systems have been slowly trained to output percepts — this is a *bottom-up* effect of perception proper, not a result of cognitive penetration or other top-down effects.

V Perspectival Sedimentation

In the mid-20th century, the French phenomenologist Maurice Merleau-Ponty, in his discussion of how a person’s bodily context conditions what they perceive, describes how our past experience creates “...a ‘world of thoughts,’ or a sediment left by our mental processes, which enables us to rely on our concepts and acquired judgements as we might on things there in front of us, presented globally, without there being any need for us to resynthesize them” ([1945] 1962: 130). To Merleau-Ponty, this notion of *sedimentation* is not merely a cognitive structuring of our perceptual outputs, but is a basic element of perception and our perceptual experiences: “My act of perception...takes advantage of work already done, of a general synthesis constituted once and for all...[which relies on]...that implicit or sedimentary body of knowledge” (238). This “sedimentary body of knowledge” can be both an individual matter of one’s personal knowledge

of their subjective context, as well as shared knowledge sedimented into cultural objects or institutions (348). Ultimately, for Merleau-Ponty, sedimentation plays a fundamental role in navigating both our physical and social environments as it defines much of how we understand our contexts and constitutes many of our habitual responses.⁸⁴

The model of perceptual conservatism proposed and defended by Susanna Siegel is one way of thinking about how sedimentation might function: a person's overall outlook, influenced both by their particular surroundings and experiences, as well as their position within a shared social construct, serves to underwrite our perceptual experiences. When that outlook erroneously alters the phenomenology of our experiences, said experiences are *hijacked* and any beliefs based on those experiences are epistemically downgraded. As described in Section II, perceptual conservatism's emphasis on phenomenological character can be contrasted with perceptual systems theory — a model of perception that treats phenomenology as, at most, an unnecessary experiential bonus to the otherwise epistemic outputs of perceptual systems (percepts).

As discussed in Section III, perceptual conservatism's phenomenological foundation ultimately struggles to explain manifest elements of our experience — namely, it cannot answer why familiar illusions are recalcitrant. Against this, perceptual system theory offers a natural explanation for the persistence of known illusions: the illusory object provokes the same output percepts, regardless of an agent's higher-order mental states (like beliefs about the illusory object).

Nevertheless, the utility of something like perceptual conservatism for explaining the effects of our worldviews on our perception is important. In Section IV, I buttressed perceptual systems theory by incorporating into it the dichotomous model of mental representations proposed and defended by Elisabeth Camp. If both logical concepts and associative

⁸⁴ For more see Kozyreva (2018) and Prinz (2018).

characterizations function as basic-level representations for higher-level thinking, then perceptual systems could output percepts with *either or both* conceptual or characterized contents. Insofar as the outputs of perception are characterized, then they are diachronically (and reciprocally) influenced by our perspectives in a manner that affects perceptual belief–fixation — and perhaps perceptual experience — without affecting the *contents* of perceptual experience.

Consequently, Camp’s perspectival epistemology — and the mental characterizations that underwrite it — offers a different method of explaining the process of sedimentation. Rather than directly altering the contents of our experience, sedimentation seeks to explain how our perspective serves to construct our understanding of those contents; for example, Merleau-Ponty points out that an individual’s habits come to bear on their agency, saying, “Here once more we must recognize a sort of sedimentation of our life: an attitude towards the world, when it has received frequent confirmation, acquires a favored status for us” ([1945] 1962: 441). When Camp says that the grip of a familiar perspective confers confidence in “knowing our way about” a context, she seems to be talking about basically the same thing (Forthcoming: §6.2).

Altogether, then, this thesis has argued that an architectural approach to understanding the perception/cognition border answers the same questions motivating Siegel’s phenomenological approach without exposing itself to the same problems. Contra Kant, thoughts without concepts might be empty, but intuitions without concepts do *not* have to be blind if we can instead use characterizations to see.

References

- Adler, Ben. 2019. “No Criminal Charges For Sacramento Police Who Shot Stephon Clark.” *Weekend Edition: Sunday*, National Public Radio. March 3, 2019. <https://www.npr.org/2019/03/03/699797212/no-criminal-charges-for-sacramento-police-who-shot-stephon-clark>
- Block, Ned. 1995. “On a Confusion About a Function of Consciousness.” *Behavioral and Brain Sciences* 18 (2): 227–287. DOI: [10.1017/s0140525x00038188](https://doi.org/10.1017/s0140525x00038188).
- Block, Ned. 1987. “Functional Role and Truth Conditions.” *Proceedings of the Aristotelian Society, Supplementary Volumes* 61: 157–181+183. <https://www.jstor.org/stable/4106836>.
- Bonjour, Laurence. 1978. “Can Empirical Knowledge Have a Foundation?” *American Philosophical Quarterly* 15 (1): 1–13. <https://www.jstor.org/stable/20009690>.
- Brandom, Robert. 2015. *From Empiricism to Expressivism: Brandom Reads Sellars*. Cambridge: Harvard University Press.
- Brogaard, Berit and Bartek Chomanski. 2015. “Cognitive Penetrability and High-Level Properties in Perception: Unrelated Phenomena?” *Pacific Philosophical Quarterly* 96: 469–486. DOI: [10.1111/papq.12111](https://doi.org/10.1111/papq.12111).
- Brownstein, Michael. 2018. *The Implicit Mind: Cognitive Architecture, the Self, and Ethics*. New York: Oxford University Press.
- Brownstein, Michael and Alex Madva. 2012a. “The Normativity of Automaticity.” *Mind & Language* 27 (4): 410–434. DOI: [10.1111/j.1468-0017.2012.01450.x](https://doi.org/10.1111/j.1468-0017.2012.01450.x).
- Brownstein, Michael and Alex Madva. 2012b. “Ethical Automaticity.” *Philosophy of the Social Sciences* 42 (1): 68–98. DOI: [10.1177/0048393111426402](https://doi.org/10.1177/0048393111426402).
- Brownstein, Michael and Jennifer Saul. 2016a. *Implicit Bias and Philosophy, Volume 1: Metaphysics and Epistemology*. Oxford: Oxford University Press
- Brownstein, Michael and Jennifer Saul. 2016b. *Implicit Bias and Philosophy, Volume 2: Moral Responsibility, Structural Injustice, and Ethics*. Oxford: Oxford University Press
- Burge, Tyler. 2010. *Origins of Objectivity*. New York: Oxford University Press.

- Camp, Elisabeth. Forthcoming. “Perspectival Complacency, Open-Mindedness and Amelioration.” In *Open-Mindedness and Perspectives*, edited by Wayne Riggs and Nancy Snow. Oxford: Oxford University Press.
- Camp, Elisabeth. 2020. “Imaginative Frames for Scientific Inquiry: Metaphors, Telling Facts, and Just-So Stories.” In *The Scientific Imagination*, edited by Arnon Levy and Peter Godfrey-Smith, 304–336. New York: Oxford University Press.
- Camp, Elisabeth. 2019. “Perspectives and Frames in Pursuit of Ultimate Understanding.” In *Varieties of Understanding: New Perspectives from Philosophy, Psychology, and Theology*, edited by Stephen R. Grimm, 17–45. New York: Oxford University Press.
- Camp, Elisabeth. 2018a. “Why Maps are Not Propositional.” In *Non-Propositional Intentionality*, edited by Alex Grzankowski and Michelle Montague, 19–45. Oxford: Oxford University Press.
- Camp, Elisabeth. 2018b. “A Dual Act Analysis of Slurs.” In *Bad Words: Philosophical Perspectives on Slurs*, edited by David Sosa, 29–59. Oxford: Oxford University Press.
- Camp, Elisabeth. 2015a. “Logical Concepts and Associative Characterizations.” In *The Conceptual Mind: New Directions in the Study of Concepts*, edited by Eric Margolis and Stephen Laurence, 591–621. Cambridge: The MIT Press.
- Camp, Elisabeth. 2015b. “Thinking with Maps.” *Philosophical Perspectives* 21 (1): 145–182. DOI: [10.1111/j.1520-8583.2007.00124.x](https://doi.org/10.1111/j.1520-8583.2007.00124.x).
- Camp, Elisabeth. 2004. “The Generality Constraint and Categorical Restrictions.” *The Philosophical Quarterly* 54 (215): 209–231. DOI: [10.1111/j.0031-8094.2004.00348.x](https://doi.org/10.1111/j.0031-8094.2004.00348.x).
- Cappelen, Herman. 2007. “Propositional Skeletons and Disquotational Reports.” *Proceedings of the Aristotelian Society* 107 (1p2): 207–227. DOI: [10.1111/j.1467-9264.2007.00219.x](https://doi.org/10.1111/j.1467-9264.2007.00219.x).
- Carey, Susan. 2011. “Concept Innateness, Concept Continuity, and Bootstrapping.” *Behavioral and Brain Sciences* 34 (3): 152–162. DOI: <https://www.doi.org/10.1017/S0140525X10003092>.
- Carey, Susan. 2009. *The Origin of Concepts*. New York: Oxford University Press.
- Chalmers, Alan. 2008. “Atom and Aether in Nineteenth-century Physical Science.” *Foundations of Chemistry* 10: 157–166. DOI: [10.1007/s10698-008-9051-8](https://doi.org/10.1007/s10698-008-9051-8).

- Churchland, Paul. 1988. "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." *Philosophy of Science* 55 (June): 167–187. DOI: [10.1086/289425](https://doi.org/10.1086/289425).
- Churchland, Paul. 1979. *Scientific Realism and The Plasticity of Mind*. New York: Cambridge University Press.
- Connolly, Kevin. 2019. *Perceptual Learning: The Flexibility of the Senses*. New York: Oxford University Press.
- Connolly, Kevin. 2014. "Perceptual Learning and the Contents of Perception." *Erkenntnis* 79: 1407–1418. DOI: [10.1007/s10670-014-9608-y](https://doi.org/10.1007/s10670-014-9608-y).
- Davidson, Donald. 2004. "Paradoxes of Rationality." In *Problems of Rationality*. Oxford: Oxford University Press.
- Devitt, Michael. 1990. "Meanings Just Ain't in the Head." In *Meaning and Method: Essays in Honor of Hilary Putnam*, edited by George Boolos, 79–104. Cambridge: Cambridge University Press.
- Evans, Gareth. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.
- Fleming, Kevin K., Carole L. Bandy, and Matthew O. Kimble. 2008. "Decisions to shoot in a weapon identification task: The influence of cultural stereotypes and perceived threat on false positive errors." *Social Neuroscience* 5 (2): 201–220. DOI: [10.1080/17470910903268931](https://doi.org/10.1080/17470910903268931).
- Fodor, Jerry. 1988. "A Reply to Churchland's "Perceptual Plasticity and Theoretical Neutrality"." *Philosophy of Science* 55 (June): 188–198. DOI: [10.1086/289426](https://doi.org/10.1086/289426).
- Fodor, Jerry. 1984. "Observation Reconsidered." *Philosophy of Science* 51 (March): 23–43. DOI: [10.1086/289162](https://doi.org/10.1086/289162).
- Fodor, Jerry. 1983. *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge: Bradford Books.
- Gendler, Tamar Szabó. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* 156 (1): 33–63. DOI: [10.1007/s11098-011-9801-7](https://doi.org/10.1007/s11098-011-9801-7).

- Gendler, Tamar Szabó. 2008a. "Alief in Action (and Reaction)." *Mind & Language* 23 (5): 552–585. DOI: [10.1111/j.1468-0017.2008.00352.x](https://doi.org/10.1111/j.1468-0017.2008.00352.x).
- Gendler, Tamar Szabó. 2008b. "Alief and Belief." *The Journal of Philosophy* 105 (10): 634–663. DOI: [10.5840/jphil20081051025](https://doi.org/10.5840/jphil20081051025).
- Ghijzen, Harmen. 2018. "How to Explain the Rationality of Perception." *Analysis* 78 (3): 500–512. DOI: [10.1093/analys/any038](https://doi.org/10.1093/analys/any038).
- Gibson, James. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin Harcourt.
- Glüer, Kathrin. 2018. "Defeating Looks." *Synthese* 195: 2985–3012. DOI: [10.1007/s11229-016-1186-x](https://doi.org/10.1007/s11229-016-1186-x)
- Glüer, Kathrin. 2009. "In Defence of a Doxastic Account of Experience." *Mind & Language* 24 (3): 297–327. DOI: [10.1111/j.1468-0017.2009.01364.x](https://doi.org/10.1111/j.1468-0017.2009.01364.x).
- Goldman, Alvin. 1967. "A Causal Theory of Knowing." *Journal of Philosophy* 64 (12): 357–372. DOI: [10.2307/2024268](https://doi.org/10.2307/2024268).
- Green, E.J. 2018. "What do Object Files Pick Out?" *Philosophy of Science* 85 (2): 177–200. DOI: [10.1086/696376](https://doi.org/10.1086/696376).
- Green, E.J. and Jake Quilty-Dunn. 2017. "What is an Object File?" *British Journal for the Philosophy of Science*, accepted manuscript 'axx055.' DOI: [10.1093/bjps/axx055](https://doi.org/10.1093/bjps/axx055).
- Grice, H.P. and Alan White. 1961. "The Causal Theory of Perception." *Aristotelian Society Supplementary Volume* 35 (1): 121–168. DOI: [10.1093/aristoteliansupp/35.1.121](https://doi.org/10.1093/aristoteliansupp/35.1.121).
- Gupta, Anil. 2006. "Experience and Knowledge." In *Perceptual Experience*, edited by Tamar Szabó Gendler and John Hawthorne, 181–204. Oxford: Oxford University Press.
- Heinrich, Joseph, Steven Heine, and Ara Norenzayan. 2010. "The Weirdest People in the World?" *Behavioral and Brain Sciences* 33 (2–3): 61–83. DOI: [10.1017/S0140525X0999152X](https://doi.org/10.1017/S0140525X0999152X).
- Johnston, Mark. 2014. "The Problem with the Content View." In *Does Perception Have Content?*, edited by Berit Brogaard, 105–137. Oxford: Oxford University Press.

- Johnston, Mark. 2006. "Better Than Mere Knowledge? The Function of Sensory Awareness." In *Perceptual Experience*, edited by Tamar Szabó Gendler and John Hawthorne, 260–290. Oxford: Oxford University Press.
- Jones, Christopher and Russell Fazio. 2010. "Person Categorization and Automatic Racial Stereotyping Effects on Weapon Identification." *Personality and Social Psychology Bulletin* 36 (8): 1073–1085. DOI: [10.1177/0146167210375817](https://doi.org/10.1177/0146167210375817).
- Kant, Immanuel. (1781) 1998. *Critique of Pure Reason*, translated and edited by Paul Guyer and Allen W. Wood. New York: Cambridge University Press.
- Korcz, Keith Allan. 2019. "Inference and the Basing Relation" in *Well-Founded Belief: New Essays on the Epistemic Basing Relation*, edited by J. Adam Carter and Patrick Bondy, 113–124. New York: Routledge.
- Kornblith, Hilary. 2015. *Second Thoughts and the Epistemological Enterprise*. New York: Cambridge University Press.
- Kozyreva, Anastasia. 2018. "Non-representational Approaches to the Unconscious in the Phenomenology of Husserl and Merleau-Ponty." *Phenomenology and the Cognitive Sciences* 17: 199–224. DOI: [10.1007/s11097-016-9492-9](https://doi.org/10.1007/s11097-016-9492-9).
- Kriegel, Uriah. 2012. "Moral Motivation, Moral Phenomenology, and the Alief/Belief Distinction." *The Australasian Journal of Philosophy* 90 (3): 469–486. DOI: [10.1080/00048402.2011.607464](https://doi.org/10.1080/00048402.2011.607464).
- Kubota, Jennifer and Tiffany Ito. 2014. "The Role of Expression and Race in Weapons Identification." *Emotion* 14 (6): 1115–1124. DOI: [10.1037/a0038214](https://doi.org/10.1037/a0038214).
- Lyons, Jack. 2016a. "Inferentialism and Cognitive Penetration of Perception." *Episteme* 13: 1–28. DOI: [10.1017/epi.2015.60](https://doi.org/10.1017/epi.2015.60).
- Lyons, Jack. 2016b. "Experiential Evidence?" *Philosophical Studies* 173: 1053–1079. DOI: [10.1007/s11098-015-0540-z](https://doi.org/10.1007/s11098-015-0540-z).
- Lyons, Jack. 2015a. "Unencapsulated Modules and Perceptual Judgment." In *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, 103–122. Oxford: Oxford University Press.

- Lyons, Jack. 2015b. "Seemings and Justification." *Analysis* 75 (1): 153–164. DOI: [10.1093/analys/anu100](https://doi.org/10.1093/analys/anu100).
- Lyons, Jack. 2011. "Circularity, Reliability, and the Cognitive Penetration of Perception." *Philosophical Issues* 21: 289–311. DOI: [10.1111/j.1533-6077.2011.00205.x](https://doi.org/10.1111/j.1533-6077.2011.00205.x).
- Lyons, Jack. 2009. *Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World*. New York: Oxford University Press.
- Lyons, Jack. 2009. "Perception and Virtue Reliabilism." *Acta Analytica* 24: 249–261. DOI: [10.1007/s12136-009-0064-2](https://doi.org/10.1007/s12136-009-0064-2).
- MacPherson, Fiona. 2012. "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism." *Philosophy and Phenomenological Research* 84: 24–62. DOI: [10.1111/j.1933-1592.2010.00481.x](https://doi.org/10.1111/j.1933-1592.2010.00481.x).
- Mandelbaum, Eric. 2018. "Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception." *Philosophy and Phenomenological Research* 97: 267–283. DOI: [10.1111/phpr.12368](https://doi.org/10.1111/phpr.12368).
- Mandelbaum, Eric. 2016. "Attitude, Inference, Association: On the Propositional Structure of Implicit Bias." *Noûs* 50 (3): 629–658. DOI: [10.1111/nous.12089](https://doi.org/10.1111/nous.12089).
- Mandelbaum, Eric. 2011. "What is the Narrow Content of Fence (and other Definitionally and Interpretationally Primitive Concepts)?" *Behavioral and Brain Science*, 34 (3): 138. DOI: [10.1017/S0140525X10002360](https://doi.org/10.1017/S0140525X10002360).
- Marr, David. 1982. *Vision: A Computational Investigation into Human Representation and Processing of Visual Information*. Cambridge: The MIT Press.
- Merleau-Ponty, Maurice. (1945) 1962. *Phenomenology of Perception*, translated by Colin Smith. London: Routledge & Kegan Paul, Ltd.
- Millikan, Ruth. 2017. *Beyond Concepts: Unicepts, Language, and Natural Information*. Oxford: Oxford University Press.
- Millikan, Ruth. 2004. *Varieties of Meaning: The 2002 Jean Nicod Lectures*. Cambridge: Bradford Books.

- Payne, B. Keith. 2006. "Weapon Bias: Split-Second Decisions and Unintended Stereotyping." *Current Directions in Psychological Science* 15 (6): 287–291. DOI: [10.1111/j.1467-8721.2006.00454.x](https://doi.org/10.1111/j.1467-8721.2006.00454.x).
- Potter, Mary, Brad Wyble, Carl Erick Haggmann, and Emily McCourt. 2014. "Detecting meaning in RSVP at 13 ms per picture." *Attention, Perception, & Psychophysics* 76 (2): 270–279. DOI: [10.3758/s13414-013-0605-z](https://doi.org/10.3758/s13414-013-0605-z).
- Prinz, Jesse. 2018. "Moral Sedimentation." In *Neuroexistentialism: Meaning, Morals, and Purpose in the Age of Neuroscience*, edited by Gregg D. Caruso and Owen Flanagan, 87–107. New York: Oxford University Press.
- Proctor, Darby and Sarah Brosnan. 2013. "Visual Processing Speed in Capuchin Monkeys (*Cebus apella*) and Rhesus Macaques (*Macaca mulatta*)." *International Journal of Comparative Psychology* 26 (2): 166–175. <https://escholarship.org/uc/item/6hg2n4qf>.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning.'" In *Mind, Language, and Reality: Philosophical Papers, Volume 2*, 215–271. Cambridge: Cambridge University Press.
- Pylyshyn, Zenon. 2003. *Seeing and Visualizing: It's Not What You Think*. Cambridge: Bradford Books.
- Pylyshyn, Zenon. 1999. "Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception." *Behavioral and Brain Sciences* 22 (3): 341–365. DOI: [10.1017/s0140525x99002022](https://doi.org/10.1017/s0140525x99002022).
- Quilty-Dunn, Jake. Forthcoming. "Perceptual Pluralism." *Noûs*. DOI: [10.1111/nous.12285](https://doi.org/10.1111/nous.12285).
- Raftopoulos, Athanassios. 2001. "Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception." *Cognitive Science* 25: 423–451. DOI: [10.1207/s15516709cog2503_4](https://doi.org/10.1207/s15516709cog2503_4).
- Rosch, Eleanor. 1999. "Reclaiming Concepts." *Journal of Consciousness Studies* 6 (11–12): 61–77.
- Rosch, Eleanor. 1978. "Principles of Categorization." In *Cognition and Categorization*, edited by Eleanor Rosch and Barbara Lloyd, 27–48. Hillsdale: L. Erlbaum Associates.

- Rosch, Eleanor, Carolyn Mervis, Wayne Gray, David Johnson, and Penny Boyes-Braem. 1976. "Basic Objects in Natural Categories." *Cognitive Psychology* 8 (3): 382–439. DOI: [10.1016/0010-0285\(76\)90013-X](https://doi.org/10.1016/0010-0285(76)90013-X).
- Sellars, Wilfrid. 1968. *Science and Metaphysics: Variation on Kantian Themes*. London: Routledge & Kegan Paul.
- Sellars, Wilfrid. (1956) 1997. *Empiricism and the Philosophy of Mind*. Cambridge: Harvard University Press.
- Schellenberg, Susanna. 2018. *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press.
- Siegel, Susanna. 2017. *The Rationality of Perception*. Oxford: Oxford University Press.
- Siegel, Susanna. 2016. "How Is Wishful Seeing Like Wishful Thinking?" *Philosophy and Phenomenological Research* 95 (2): 408–435. DOI: [10.1111/phpr.12273](https://doi.org/10.1111/phpr.12273).
- Siegel, Susanna. 2012. "Cognitive Penetrability and Perceptual Justification" *Noûs* 46 (2): 201–222. DOI: [10.1111/j.1468-0068.2010.00786.x](https://doi.org/10.1111/j.1468-0068.2010.00786.x).
- Siegel, Susanna. 2011. *The Contents of Visual Experience*. New York: Oxford University Press.
- Stokes, Dustin. 2018. "Rich Perceptual Content and Aesthetic Properties." In *Evaluative Perception*, edited by Anna Bergqvist and Robert Cowan, 19–41. Oxford: Oxford University Press.
- Tucker, Chris. 2013. "Seeming and Justification: An Introduction." In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, edited by Chris Tucker, 1–29. Oxford: Oxford University Press.
- Turri, John. 2019. "Reasons and Basing in Commonsense Epistemology: Evidence from Two Experiments." in *Well-Founded Belief: New Essays on the Epistemic Basing Relation*, edited by J. Adam Carter and Patrick Bondy, 92–112. New York: Routledge.
- Turri, John. 2011. "Believing for a Reason." *Erkenntnis* 74: 383–397. DOI: [10.1007/s10670-011-9271-5](https://doi.org/10.1007/s10670-011-9271-5).

- Väyrynen, Pekka. 2019. "Thick Ethical Concepts." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2019/entries/thick-ethical-concepts/>.
- Vigdor, Neil. 2019. "Stephon Clark's Sons Reach \$2.4 Million Settlement Over Police Killing." *New York Times*. October 10, 2019. <https://www.nytimes.com/2019/10/10/us/stephon-clark-shooting-settlement.html>
- Wamsley, Laurel. 2018. "Video Shows Sacramento Police Shooting Unarmed Black Man In Grandparents' Backyard." *National Public Radio*. March 22, 2018. <https://www.npr.org/sections/thetwo-way/2018/03/22/596051907/video-shows-sacramento-police-shoot-unarmed-black-man-in-grandparents-backyard>
- Weiskrantz, Lawrence. 1986. *Blindsight: A Case Study and Implications*. Oxford: Oxford University Press.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations*, translated by G. E. M. Anscombe. Oxford: Basil Blackwell.